



US005477451A

United States Patent [19]

Brown et al.

[11] Patent Number: 5,477,451
[45] Date of Patent: Dec. 19, 1995

[54] METHOD AND SYSTEM FOR NATURAL LANGUAGE TRANSLATION

[75] Inventors: Peter F. Brown, New York; John Cocke, Bedford; Stephen A. Della Pietra, Pearl River; Vincent J. Della Pietra, Blauvelt; Frederick Jelinek, Briarcliff Manor; Jennifer C. Lai, Garrison; Robert L. Mercer, Yorktown Heights, all of N.Y.

[73] Assignee: International Business Machines Corp., Yorktown Heights, N.Y.

[21] Appl. No.: 736,278

[22] Filed: Jul. 25, 1991

[51] Int. Cl.⁶ G06F 17/20; G06F 17/27

[52] U.S. Cl. 364/419.08

[58] Field of Search 364/419, 419.02, 364/419.08, 419.16; 381/43, 51

[56] References Cited

U.S. PATENT DOCUMENTS

4,754,489	6/1988	Bokser	382/40
4,829,580	5/1989	Church	381/52
4,852,173	7/1989	Bahl et al.	381/43
4,882,759	11/1989	Bahl et al.	381/51
4,984,178	1/1991	Hemphill et al.	364/419
4,991,094	2/1991	Fagen et al.	364/419
5,033,087	7/1991	Bohl et al.	381/43
5,068,789	11/1991	Van Vliembergen	364/419
5,109,509	4/1992	Katayama et al.	395/600
5,146,405	9/1992	Church	364/419
5,200,893	4/1993	Ozawa et al.	364/419

FOREIGN PATENT DOCUMENTS

0327266	8/1989	European Pat. Off. .
0357344	3/1990	European Pat. Off. .
0399533	11/1990	European Pat. Off. .
WO9010911	9/1990	WIPO .

OTHER PUBLICATIONS

"Method for Inferring Lexical Associations from Textual Co-Occurrences", IBM Technical Disclosure Bulletin, vol. 33 1B, Jun. 1990.

"Tagging Text with a Probabilistic Model", by Bernard Merialdo, *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, Paris, France, May 14-17, 1991.

(List continued on next page.)

Primary Examiner—David M. Huntley

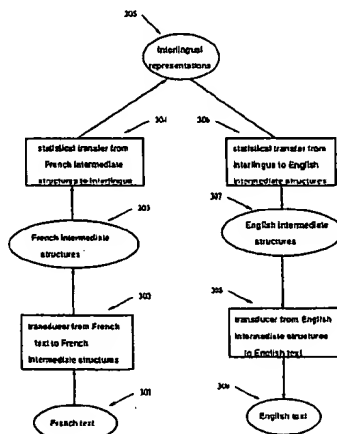
Assistant Examiner—Charles R. Kyle

Attorney, Agent, or Firm—Sterne, Kessler, Goldstein & Fox; Robert P. Tassinari

[57] ABSTRACT

The present invention is a system for translating text from a first source language into a second target language. The system assigns probabilities or scores to various target-language translations and then displays or makes otherwise available the highest scoring translations. The source text is first transduced into one or more intermediate structural representations. From these intermediate source structures a set of intermediate target-structure hypotheses is generated. These hypotheses are scored by two different models: a language model which assigns a probability or score to an intermediate target structure, and a translation model which assigns a probability or score to the event that an intermediate target structure is translated into an intermediate source structure. Scores from the translation model and language model are combined into a combined score for each intermediate target-structure hypothesis. Finally, a set of target-text hypotheses is produced by transducing the highest scoring target-structure hypotheses into portions of text in the target language. The system can either run in batch mode, in which case it translates source-language text into a target language without human assistance, or it can function as an aid to a human translator. When functioning as an aid to a human translator, the human may simply select from the various translation hypotheses provided by the system, or he may optionally provide hints or constraints on how to perform one or more of the stages of source transduction, hypothesis generation and target transduction.

42 Claims, 54 Drawing Sheets



OTHER PUBLICATIONS

- "Word-Sense Disambiguation Using Statistical Methods", by Peter F. Brown et al., appearing in the *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Jun. 1991, pp. 264-270.
- "Lex—A Lexical Analyzer Generator", M. E. Lesk, *Computer Science Technical Report*, No. 39, Bell Laboratories, Oct. 1975.
- "Self Organized Language Modeling for Speech Recognition" by F. Jelinek, *Language Processing for Speech Recognition*, pp. 450-506.
- "A Tree-Based Statistical Language Model for Natural Language Speech Recognition" by Lalit R. Bahl et al., *IEEE Transactions on Acoustics*, vol. 37, No. 7, Jul. 1989, pp. 1001-1008.
- "Trainable Grammars for Speech Recognition" by James K. Baker, *Speech Communications Papers Presented at the 97th Meeting of the Acoustic Society of America*, 1979, pp. 547-550.
- "Interpolated Estimation of Markov Source Parameters from Sparse Data" by F. Jelinek and R. L. Mercer, *Workshop on Pattern Recognition in Practice*, Amsterdam (Netherlands): North Holland, May 21-23, 1980.
- "An Inequality and Associated Maximization Technique in Statistical Estimation for Probabilistic Functions of Markov Processes", by Leonard E. Baum, *Inequalities*, vol. 3, 1972, pp. 1-8.
- "Aligning Sentences in Parallel Corpora" by Peter F. Brown et al., appearing in the *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Jun. 1991, pp. 169-176.
- "Partial Traceback in Continuous Speech Recognition" by James C. Spohrer et al., *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Paris, France, 1982.
- "Deriving Translation Data from Bilingual Texts", Catizone et al., *Proceedings of the First International Acquisition Workshop*, Detroit, Mich., 1989.
- "Making Connections", by M. Kay, appearing in *ACH/ALLC '91*, Tempe, Ariz., 1991, p. 1.

US-PAT-NO: 5477451

DOCUMENT-IDENTIFIER: US 5477451 A

****See image for Certificate of Correction****

TITLE: Method and system for natural language translation

Abstract Text - ABTX (1):

The present invention is a system for translating text from a first source language into a second target language. The system assigns probabilities or scores to various target-language translations and then displays or makes otherwise available the highest scoring translations. The source text is first transduced into one or more intermediate structural representations. From these intermediate source structures a set of intermediate target-structure hypotheses is generated. These hypotheses are scored by two different models: a language model which assigns a probability or score to an intermediate target structure, and a translation model which assigns a probability or score to the event that an intermediate target structure is translated into an intermediate source structure. Scores from the translation model and language model are combined into a combined score for each intermediate target-structure hypothesis. Finally, a set of target-text hypotheses is produced by transducing the highest scoring target-structure hypotheses into portions of text in the target language. The system can either run in batch mode, in which case it translates source-language text into a target language without human assistance, or it can function as an aid to a human translator. When functioning as an aid to a human translator, the human may simply select from the various translation hypotheses provided by the system, or he may optionally provide hints or constraints on how to perform one or more of the stages of source transduction, hypothesis generation and target transduction.

Brief Summary Text - BSTX (6):

This impasse was overcome in the early 1970's with the introduction of statistical techniques to speech recognition. In the statistical approach, linguistic rules are extracted automatically using statistical techniques from large databases of speech and text. Different types of linguistic information are combined via the formal laws of probability theory. Today, almost all speech recognition systems are based on statistical techniques.

Brief Summary Text - BSTX (7):

Speech recognition has benefited by using statistical language models which exploit the fact that not all word sequences occur naturally with equal probability. One simple model is the trigram model of English, in which it is assumed that the probability that a word will be spoken depends only on the previous two words that have been spoken. Although trigram models are simple-minded, they have proven extremely powerful in their ability to predict words as they occur in natural language, and in their ability to improve the performance of natural-language speech recognition. In recent years more sophisticated language models based on probabilistic decision-trees, stochastic context-free grammars and automatically discovered classes of words have also been used.

Brief Summary Text - BSTX (8):

In the early days of speech recognition, acoustic models were created by linguistic experts, who expressed their knowledge of acoustic-phonetic rules in programs which analyzed an input speech signal and produced as output a sequence of phonemes. It was thought to be a simple matter to decode a word sequence from a sequence of phonemes. It turns out, however, to be a very difficult job to determine an accurate phoneme sequence from a speech signal. Although human experts certainly do exist, it has proven extremely difficult to formalize their knowledge. In the alternative statistical approach, statistical models, most predominantly hidden Markov models, capable of learning acoustic-phonetic knowledge from samples of speech are employed.

Brief Summary Text - BSTX (10):

Statistical techniques in speech recognition provide two advantages over the rule-based approach. First, they

provide means for automatically extracting information from large bodies of acoustic and textural data, and second, they provide, via the formal rules of probability theory, a systematic way of combining information acquired from different sources. The problem of machine translation between natural languages is an entirely different problem than that of speech recognition. In particular, the main area of research in speech recognition, acoustic modeling, has no place in machine translation. Machine translation does face the difficult problem of coping with the complexities of natural language. It is natural to wonder whether this problem won't also yield to an attack by statistical methods, much as the problem of coping with the complexities of natural speech has been yielding to such an attack. Although the statistical models needed would be of a very different nature, the principles of acquiring rules automatically and combining them in a mathematically principled fashion might apply as well to machine translation as they have to speech recognition.

Brief Summary Text - BSTX (14):

A target-structure language model is used to estimate a first score which is proportional to the probability of occurrence of each intermediate target-structure of text associated with the target hypotheses. A target structure-to-source-structure translation model is used to estimate a second score which is proportional to the probability that the intermediate target-structure of text associated with the target hypotheses will translate into the intermediate source-structure of text. For each target hypothesis, the first and second scores are combined to produce a target hypothesis match score.

Brief Summary Text - BSTX (18):

The intermediate target structures may be expressed as an ordered sequence of morphological units, and the first score probability may be obtained by multiplying the conditional probabilities of each morphological unit within an intermediate target structure given the occurrence of previous morphological units within the intermediate target structure. In another embodiment, the conditional probability of each unit of each of the intermediate target structure may depend only on a fixed number of preceding units within the intermediate target structure.

Drawing Description Text - DRTX (66):

FIG. 64 is an example of a subset lattice;

Detailed Description Text - DETX (104):

1. A language model 101 which assigns a probability or score $P(E)$ to any portion of English text E;

Detailed Description Text - DETX (105):

2. A translation model 102 which assigns a conditional probability or score $P(F|E)$ to any portion of French text F given any portion of English text E; and

Detailed Description Text - DETX (106):

3. A decoder 103 which given a portion of French text F finds a number of portions of English text E, each of which has a large combined probability or score

Detailed Description Text - DETX (108):

A shortcoming of the simplified architecture depicted in FIG. 1 is that the language model 101 and the translation model 102 deal directly with unanalyzed text. The linguistic information in a portion of text and the relationships between translated portions of text is complex, involving linguistic phenomena of a global nature. However, the models 101 and 102 must be relatively simple so that their parameters can be reliably estimated from a manageable amount of training data. In particular, they are restricted to the description of local structure.

Detailed Description Text - DETX (115):

4. an English structure language model 204 which assigns a probability or score $P(E')$ to any intermediate structure E' ;

Detailed Description Text - DETX (116):

5. an English structure to French structure translation model 205 which assigns a conditional probability or score $P(F'.\text{vertline}.E')$ to any intermediate structure F' given any intermediate structure E' ; and

Detailed Description Text - DETX (117):

6. a decoder 206 which given a French structure F' finds a number of English structures E' , each of which has a large combined probability or score

Detailed Description Text - DETX (127):

In step 505 in FIG. 5, the user-aided system, portions of source text may be selected by the user. The user might, for example, select a whole document, a sentence, or a single word. The system might then show the users possible translations of the selected portion of text. For example, if the user selected only a single word the system might show a ranked list of possible translations of that word. The ranks being determined by statistical models that would be used to estimate the probabilities that the source word translates in various manners in the source context in which the source word appears.

Detailed Description Text - DETX (130):

FIG. 7 depicts in more detail the step 404 of the batch translation system 401 depicted in FIG. 4. The step 404 is expanded into four steps. In the first step 701 the input source text 707 is transduced to one or more intermediate source structures. In the second step 702 a set of one or more hypothesized target structures are generated. This step 702 makes use of a language model 705 which assigns probabilities or scores to target structures and a translation model 706 which assigns probabilities or scores to source structures given target structures. In the third step 703 the highest scoring hypothesis is selected. In the fourth step 704 the hypothesis selected in step 703 is transduced into text in the target language 708.

Detailed Description Text - DETX (139):

The computer platform 1014 typically includes an operating system 1003. A data storage device 1002 is also called a secondary storage and may include hard disks and/or tape drives and their equivalents. The data storage device 1002 represents non-volatile storage. The data storage 1002 may be used to store data for the language and translation models components of the translation system 1001.

Detailed Description Text - DETX (140):

Various peripheral components may be connected to the computer platform 1014, such as a terminal 1012, a microphone 1008, a keyboard 1013, a scanning device 1009, an external network 1010, and a printing device 1011. A user 503 may interact with the translation system 1001 using the terminal 1012 and the keyboard 1013, or the microphone 1008, for example. As another example, the user 503 might receive a document from the external network 1010, translate it into another language using the translation system 1001, and then send the translation out on the external network 1010.

Detailed Description Text - DETX (142):

The translation system can receive source text in a variety of known manners. The following are only a few examples of how the translation system receives source text (e.g. data), to be translated. Source text to be translated may be directly entered into the computer system via the keyboard 1013. Alternatively, the source text could be scanned in using the scanner 1009. The scanned data could be passed through a character recognizer in a known manner and then passed on to the translation system for translation. Alternatively, the user 503 could identify the location of the source text in main or secondary storage, or perhaps on removable secondary storage (such as on a floppy disk), the computer system could retrieve and then translate the text accordingly. As a final

example, with the addition of a speech recognition component, it would also be possible to speak into the microphone 1008, have the speech converted into source text by the speech recognition component.

Detailed Description Text - DETX (143):

Translated target text produced by the translation application running on the computer system may be output by the system in a variety of known manners. For example, it may be displayed on the terminal 1012, stored in RAM 1005, stored data storage 1002, printed on the printer 1011, or perhaps transmitted out over the external network 1010. With the addition of a speech synthesizer it would also be possible to convert translated target text into speech in target language.

Detailed Description Text - DETX (144):

Step 403 in FIG. 4 and step 505 in FIG. 5 measure, receive or otherwise capture a portion of source text to be translated. In the context of this invention, text is used to refer to sequences of characters, formatting codes, the typographical marks. It can be provided to the system in a number of different fashions, such as on a magnetic disk, via a computer network, as the output of an optical scanner, or of a speech recognition system. In some preferred embodiments, the source text is captured a sentence at a time. Source text is parsed into sentences using a finite-state machine which examines the text for such things as upper and lower case characters and sentence terminal punctuations. Such a machine can easily be constructed by someone skilled in the art. In other embodiments, text may be parsed into units such as phrases or paragraphs which are either smaller or larger than individual sentences.

Detailed Description Text - DETX (155):

It should be understood that FIG. 11 represents only one embodiment of the source-transducer 701. Many variations are possible. For example, the transducers 1101, 1102, 1103, 1104, 1105, 1106 may be augmented and/or replaced by other transducers. Other embodiments of the source-transducer 701 may include a transducer that groups words into compound words or identifies idioms. In other embodiments, rather than a single intermediate source-structure being produced for each source sentence, a set of several intermediate source-structures together with probabilities or scores may be produced. In such embodiments the transducers depicted in FIG. 11 can be replaced by transducers which produce at each stage intermediate structures with probabilities or scores. In addition, the intermediate source-structures produced may be different. For example, the intermediate structures may be entire parse trees, or case frames for the sentence, rather than a sequence of morphological units. In these cases, there may be more than one intermediate source-structure for each sentence with scores, or there may be only a single intermediate source-structure.

Detailed Description Text - DETX (215):

Referring again to FIG. 11, the transducer 1103 annotates words with part-of-speech labels. These labels are used by the subsequent transducers depicted the figure. In some embodiments of transducer 1103, part-of-speech labels are assigned to a word sequence using a technique based on hidden Markov models. A word sequence is assigned the most probable part-of-speech sequence according to a statistical model, the parameters of which are estimated from large annotated texts and other even larger un-annotated texts. The technique is fully explained in article by Bernard Merialdo entitled 'Tagging text with a Probabilistic Model' in the Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, May 14-17, 1991. This article is incorporated by reference herein.

Detailed Description Text - DETX (327):

Registers: Just knowing that a regular expression matches an input attribute tuple sequence usually does not provide enough information for the construction of an appropriate output attribute tuple sequence. Data is usually also required about the attribute tuples matched by different elements of the regular expression. In ordinary LEX, to extract this type of information often requires the matched input sequence to be parsed again. To avoid this cumbersome approach, the pattern-matcher 1601 makes details about the positions in the input stream of the matched elements of the regular expression more directly available. From these positions, the identities of the attribute tuples can then be determined.

Detailed Description Text - DETX (354):

It should be understood that FIG. 19 represents only one possible embodiment of the target-transducer 704. Many variations are possible. For example, in other embodiments, rather than a single target sentence being produced for each intermediate target-structure, a set of several target sentences together with probabilities or scores may be produced. In such embodiments, the transducers depicted in FIG. 19 can be replaced by transducers which produce at each stage several target sentences with probabilities or scores. Moreover, in embodiments of the present invention in which the intermediate target-structures are more sophisticated than lexical morph sequences, the target-structure transducer is also more involved. For example, if the intermediate target-structure consists of a parse tree of a sentence or case frames for a sentence, then the target-structure transducer converts these to the target language.

Detailed Description Text - DETX (376):

To perform such conversions, the transducer 1909 uses a target-language model to assign a probability or score to each of the different possible sentences corresponding to an input sentence with a morphological ambiguity. The sentence with the highest probability or score is selected. In the above example, the sentence how quickly can you run? is selected because it has a higher target-language model probability or score than how quick can you run? In some embodiments of the transducer 1909 the target-language model is a trigram model similar to the target-structure language model 706. Such a transducer can be constructed by a person skilled in the art. The last transducer 1910 assigns a case to the words of a target sentence based on the casing rules for English. Principally this involves capitalizing the words at the beginning of sentences. Such a transducer can easily be constructed by a person skilled in the art.

Detailed Description Text - DETX (382):

The inventions described in this specification employ probabilistic models of the target language in a number of places. These include the target structure language model 705, and the class language model used by the decoder 404. As depicted in FIG. 20, the role of a language model is to compute an a priori probability or score of a target structure.

Detailed Description Text - DETX (388):

Since it is expensive to evaluate a language model in the context of a complete system, it is useful to have an intrinsic measure of the quality of a language model. One such measure is the probability that the model assigns to the large sample of target structures. One judges as better the language model which yields the greater probability. When the target structure is a sequence of words or morphs, this measure can be adjusted so that it takes account of the length of the structures. This leads to the notion of the perplexity of a language model with respect to a sample of text S: $2^{\frac{1}{|S|} \sum_{i=1}^{|S|} -\log_2 p(m_i)}$ where $|S|$ is the number of morphs of S. Roughly speaking, the perplexity is the average number of morphs which the model cannot distinguish between, in predicting a morph of S. The language model with the smaller perplexity will be the one which assigns the larger probability to S.

Detailed Description Text - DETX (392):

Suppose $m_{sub.1}, m_{sub.2}, m_{sub.3}, \dots, m_{sub.k}$ be a sequence of k morphs $m_{sub.i}$. For $1 \leq i \leq k$, let $m_{sub.i}^{sup.i}$ denote the subsequence $m_{sub.i}, m_{sub.i+1}, \dots, m_{sub.k}$. For any sequence, the probability of a $m_{sub.1}^{sup.k}$ is equal to the product of the conditional probabilities of each morph $m_{sub.i}$ given the previous morphs $m_{sub.1}^{sup.i-1}$: $\prod_{i=1}^k p(m_{sub.i} | m_{sub.1}^{sup.i-1})$

Detailed Description Text - DETX (394):

For an n-gram model, the conditional probability of a morph in a sequence is assumed to depend on its history only through the previous n-1 morphs:

Detailed Description Text - DETX (395):

For a vocabulary of size V , a 1-gram model is determined by $V-1$ independent numbers, one probability $\Pr(m)$ for each morph m in the vocabulary, minus one for the constraint that all of the probabilities add up to 1. A 2-gram model is determined by $V \cdot \sup{2} - 1$ independent numbers, $V(V-1)$ conditional probabilities of the form $\Pr(m.\text{sub}.2.\text{vertline}.m.\text{sub}.1)$ and $V-1$ of the form $\Pr(m)$. In general, an n -gram model is determined by $V \cdot \sup{n} - 1$ independent numbers, $V \cdot \sup{n-1} (V-1)$ conditional probabilities of the form $\Pr(m.\text{sub}.n.\text{vertline}.m.\text{sub}.1.\text{sup}.n-1)$, called the order- n conditional probabilities, plus $V \cdot \sup{n-1} - 1$ numbers which determine an $(n-1)$ gram model.

Detailed Description Text - DETX (396):

The order- n conditional probabilities of an n -gram model form the transition matrix of an associated Markov model. The states of this Markov model are sequences of $n-1$ morphs, and the probability of a transition from the state $m.\text{sub}.1\ m.\text{sub}.2\ \dots\ m.\text{sub}.n-1$ to the state $m.\text{sub}.2\ m.\text{sub}.3\ \dots\ m.\text{sub}.n$ is $\Pr(m.\text{sub}.n.\text{vertline}.m.\text{sub}.1\ m.\text{sub}.2\ \dots\ m.\text{sub}.n-1)$. An n -gram language model is called consistent if, for each string $m.\text{sub}.1.\text{sup}.n-1$, the probability that the model assigns to $m.\text{sub}.1.\text{sup}.n-1$ is the steady state probability for the state $m.\text{sub}.1.\text{sup}.n-1$ of the associated Markov model.

Detailed Description Text - DETX (398):

The simplest form of an n -gram model is obtained by assuming that all the independent conditional probabilities are independent parameters. For such a model, values for the parameters can be determined from a large sample of training text by sequential maximum likelihood training. The order n -probabilities are given by $\#\text{EQU3}\#$ where $f(m.\text{sub}.1.\text{sup}.i)$ is the number of times the string of morphs $m.\text{sub}.1.\text{sup}.i$ appears in the training text. The remaining parameters are determined inductively by an analogous formula applied to the corresponding $n-1$ -gram model. Sequential maximum likelihood training does not produce a consistent model, although for a large amount of training text, it produces a model that is very nearly consistent.

Detailed Description Text - DETX (399):

Unfortunately, many of the parameters of a simple n -gram model will not be reliably estimated by this method. The problem is illustrated in Table 3, which shows the number of 1-2-, and 3-grams appearing with various frequencies in a sample of 365,893,263 words of English text from a variety of sources. The vocabulary consists of the 260,740 different words plus a special unknown word into which all other words are mapped. Of the $6.799 \cdot \text{times} \cdot 10 \cdot \sup{10}$ 2-grams that might have occurred in the data, only 14,494,217 actually did occur and of these, 8,045,024 occurred only once each. Similarly, of the $1.733 \cdot \text{times} \cdot 10 \cdot \sup{16}$ 3-grams that might have occurred, only 75,349,888 actually did occur and of these, 53,737,350 occurred only once each. These data and Turing's formula imply that 14.7 percent of the 3-grams and for 2.2 percent of the 2-grams in a new sample of English text will not appear in the original

Detailed Description Text - DETX (400):

sample. Thus, although any 3-gram that does not appear in the original sample is rare, there are so many of them that their aggregate probability is substantial.

Detailed Description Text - DETX (403):

A solution to this difficulty is provided by interpolated estimation, which is described in detail in the paper "Interpolated estimation of Markov source parameters from sparse data", by F. Jelinek and R. Mercer and appearing in Proceeding of the Workshop on Pattern Recognition in Practice, published by North-Holland, Amsterdam, The Netherlands, in May 1980. Interpolated estimation combines several models into a smoothed model which uses the probabilities of the more accurate models where they are reliable and, where they are unreliable, falls back on the more reliable probabilities of less accurate models. If $\Pr.\text{sup}.(j)$ ($m.\text{sub}.i.\text{vertline}.m.\text{sub}.1.\text{sup}.i-1$) is the j th language model, the smoothed model, \Pr ($m.\text{sub}.i.\text{vertline}.m.\text{sub}.1.\text{sup}.i-1$), is given by $\#\text{EQU4}\#$

Detailed Description Text - DETX (404):

The values of the $\lambda.\text{sub}.j$ ($m.\text{sub}.1.\text{sup}.i-1$) are determined using the EM method, so as to maximize the

probability of some additional sample of training text called held-out data. When interpolated estimation is used to combine simple 1-, 2-, and 3-gram models, the λ 's can be chosen to depend on $m_{sub.i-1}^{sup.i-1}$ only through the count of $m_{sub.i-2}^{sup.i-1}$. Where this count is high, the simple 3-gram model will be reliable, and, where this count is low, the simple 3-gram model will be unreliable.

Detailed Description Text - DETX (405):

The inventors constructed an interpolated 3-gram model in which the λ 's were divided into 1782 different sets according to the 2-gram counts, and determined from a held-out sample of 4,630,934 million words. The power of the model was tested using the 1,014,312 word Brown corpus. This well known corpus, which contains a wide variety of English text, is described in the book Computational Analysis of Present-Day American English, by H. Kucera and W. Francis, published by Brown University Press, Providence, R.I., 1967. The Brown corpus was not included in either the training or held-out data used to construct the model. The perplexity of the interpolated model with respect to the Brown corpus was 244.

Detailed Description Text - DETX (407):

Clearly, some words are similar to other words in their meaning and syntactic function. For example, the probability distribution of words in the vicinity of Thursday is very much like that for words in the vicinity of Friday. Of course, they will not be identical: people rarely say Thank God it's Thursday! or worry about Thursday the 13^{sup.th}.

Detailed Description Text - DETX (410):

In a simple n-gram class model, the $C_{sup.n-1+V-C}$ independent probabilities are treated as independent parameters. For such a model, values for the parameters can be determined by sequential maximum likelihood training. The order n probabilities are given by $\# \# EQU5 \# \#$ where $f(c_{sub.1}^{sup.i})$ is the number of times that the sequence of classes $c_{sub.1}^{sup.i}$ appears in the training text. (More precisely, $f(c_{sub.1}^{sup.i})$ is the number of distinct occurrences in the training text of a consecutive sequence of morphs $m_{sub.1}^{sup.i}$ for which $c_{sub.k} = c(m_{sub.k})$ for $1 \leq k \leq i$.)

Detailed Description Text - DETX (415):

A general scheme for clustering a vocabulary into classes is depicted schematically in FIG. 31. It takes as input a desired number of classes C 3101, a vocabulary 3102 of size V, and a model 3103 for a probability distribution $P(w_{sub.1}, w_{sub.2})$ over bigrams from the vocabulary. It produces as output a partition 3104 of the vocabulary into C classes. In one application, the model 3103 can be a 2-gram language model as described in Section 6, in which case $P(w_{sub.1}, w_{sub.2})$ would be proportional to the number of times that the bigram $w_{sub.1} w_{sub.2}$ appears in a large corpus of training text.

Detailed Description Text - DETX (416):

Let the score $\psi(C)$ of a partition C be the average mutual information between the classes of C with respect to the probability distribution $P(w_{sub.1}, w_{sub.2})$: $\# \# EQU6 \# \#$

Detailed Description Text - DETX (433):

The implementation can improved further by keeping track of those pairs l,m for which $p_{sub.k}(l,m)$ is different from zero. For example, suppose that P is given by a dimple bigram model trained on the data described in Table 3 of Section 6. In this case, of the $6,799 \times 10^{10}$ possible word 2-grams $w_{sub.1}, w_{sub.2}$, only 14,494,217 have non-zero probability. Thus, in this case, the sums required in equation 18 have, on average, only about 56 non-zero terms instead of 260,741 as might be expected from the size of the vocabulary.

Detailed Description Text - DETX (439):

The methods described above were used divide the 260,741-word vocabulary of Table 3, Section 6, into 1000 classes. Table 4 shows some of the classes that are particularly interesting, and Table 5 shows classes that were

selected at random. Each of the lines in the tables contains members of a different class. The average class has 260 words. The table shows only those words that occur at least ten times, and only the ten most frequent words of any class. (The other two months would appear with the class of months if this limit had been extended to twelve). The degree to which the classes capture both syntactic and semantic aspects of English is quite surprising given that they were constructed from nothing more than counts of bigrams. The class [that tha theat] is interesting because although tha and theat are English words, the method has discovered that in the training data each of them is most often a mistyped that.

Detailed Description Text - DETX (446):

As illustrated in FIG. 21, a target structure to source structure translation model $P_{\text{sub}}(\theta, f|e)$ with parameters θ is a method for calculating a conditional probability, or likelihood, $P_{\text{sub}}(\theta, f|e)$, for any source structure f given any target structure e . Examples of such structures include, but are not limited to, sequences of words, sequences of linguistic morphs, parse trees, and case frames. The probabilities satisfy: $\sum_e P_{\text{sub}}(\theta, f|e) = 1$ where the sum ranges over all structures f , and failure is a special symbol. $P_{\text{sub}}(\theta, f|e)$ can be interpreted as the probability that a translator will produce f when given e , and $P_{\text{sub}}(\theta, \text{failure}|e)$ can be interpreted as the probability that he will produce no translation when given e . A model is called deficient if $P_{\text{sub}}(\theta, \text{failure}|e)$ is greater than zero for some e .

Detailed Description Text - DETX (449):

One training methodology is maximum likelihood training, in which the parameter values are chosen so as to maximize the probability that the model assigns to a training sample consisting of a large number S of translations (f_s, e_s) , $s=1, 2, \dots, S$. This is equivalent to maximizing the log likelihood objective function $\sum_s \log P_{\text{sub}}(\theta, f_s|e_s)$

Detailed Description Text - DETX (454):

In some embodiments, illustrated in FIG. 24, a translation model 706 computes the probability of a source structure given a target structure as the sum of the probabilities of all alignments between these structures:

Detailed Description Text - DETX (455):

In some embodiments, a translation model 706 can compute the probability of a source structure given a target structure as the maximum of the probabilities of all alignments between these structures:

Detailed Description Text - DETX (456):

As depicted in FIG. 25, the probability $P_{\text{sub}}(\theta, f|e)$ of a single alignment is computed by a detailed translation model 2101. The detailed translation model 2101 employs a table 2501 of values for the parameters θ .

Detailed Description Text - DETX (463):

The probability of an alignment and source structure given a target structure is obtained by combining the probabilities computed by each of these sub-models. Corresponding to these sub-models, the table of parameter values 2501 comprises:

Detailed Description Text - DETX (464):

1a. fertility probabilities $n(\phi, e)$ where ϕ is any non-negative integer and e is any target morph;

Detailed Description Text - DETX (465):

b. null fertility probabilities $n_{\text{sub}}(\phi, m')$, where ϕ is any non-negative integer and m' is any positive integer;

Detailed Description Text - DETX (466):

2a. lexical probabilities $t(f.\text{verline}.e)$, where f is any source morph, and e is any target morph;

Detailed Description Text - DETX (467):

b. lexical probabilities $t(f.\text{verline}.*\text{null}*)$, where f is any source morph, and $*\text{null}*$ is a special symbol;

Detailed Description Text - DETX (468):

3. distortion probabilities $a(j.\text{verline}.i,m)$, where m is any positive integer, i is any positive integer, and j is any positive integer between 1 and m .

Detailed Description Text - DETX (469):

This embodiment of the detailed translation model 2101 computes the probability $P.\text{sub}..\theta.$ ($f,a.\text{verline}.e$) of an alignment a and a source structure f given a target structure e as follows. If any source entry is connected to more than one target entry, then the probability is zero. Otherwise, the probability is computed by the formula

Detailed Description Text - DETX (483):

The components of the probability $P.\text{sub}..\theta.$ ($f,a.\text{verline}.e$) are ##EQU18##

Detailed Description Text - DETX (497):

which follows because the logarithm is concave. In fact, for any e and f , ##EQU20## between probability distributions p and q . However, whereas the relative entropy is never negative, R can take any value. The inequality (35) for R is the analog of the inequality $D>0$ for D .

Detailed Description Text - DETX (513):

To apply these procedures, it is necessary to solve the maximization problems of Steps 2802 and 2901. For the models described below, this can be done explicitly. To see the basic form of the solution, suppose $P.\text{sub}..\theta.$ is a simple model given by ##EQU21## where the $\theta.(\omega.).\omega..\epsilon..OMEGA.$, are real-valued parameters satisfying the constraints ##EQU22## and for each $\omega.$ and (a,f,e) , $c(\omega.;a,f,e)$ is a non-negative integer. ² It is natural to interpret $\theta.(\omega.)$ as the probability of the event $\omega.$ and $c(\omega.;a,f,e)$ as the number of times that this event occurs in (a,f,e) . Note that

Detailed Description Text - DETX (515):

These formulae can easily be generalized to models of the form (38) for which the single equality constraint in Equation (39) is replaced by multiple constraints ##EQU26## where the subsets $OMEGA.\text{sub}..\mu.$, $\mu.=1,2,\dots$ form a partition of $OMEGA.$. Only Equation (42) needs to be modified to include a different $\lambda.\text{sub}..\mu.$ for each $\mu.$; if $\omega..\epsilon..OMEGA.\text{sub}..\mu.$, then ##EQU27##

Detailed Description Text - DETX (518):

For simple models, such as Model 1 and Model 2 described below, it is possible to calculate these counts exactly by including the contribution of each possible alignments. For more sophisticated models, such as Model 3, Model 4, and Model 5 described below, the sum over alignments in Equation 44 is too costly to compute exactly. Rather, only the contributions from a subset of alignments can be practically included. If the alignments in this subset account for most of the probability of a translation, then this truncated sum can still be a good approximation.

Detailed Description Text - DETX (520):

In calculating the counts using the update formulae 44, approximate the sum by including only the contributions from some subset of alignments of high probability.

Detailed Description Text - DETX (521):

If only the contribution from the single most probable alignment is included, the resulting procedure is called Viterbi Parameter Estimation. The most probable alignment between a target structure and a source structure is called the Viterbi alignment. The convergence of Viterbi Estimation is easily demonstrated. At each iteration, the parameter values are re-estimated so as to make the current set of Viterbi alignments as probable as possible; when these parameters are used to compute a new set of Viterbi alignments, either the old set is recovered or a set which is yet more probable is found. Since the probability can never be greater than one, this process surely converge. In fact, it converge in a finite, though very large, number of steps because there are only a finite number of possible alignments for any particular translation.

Detailed Description Text - DETX (525):

Model 1 is very simple but it is useful because its likelihood function is concave and consequently has a global maximum which can be found by the EM procedure. Model 2 is a slight generalization of Model 1. For both Model 1 and Model 2, the sum over alignments for the objective function and the relative objective function can be computed very efficiently. This significantly reduces the computational complexity of training. Model 3 is more complicated and is designed to more accurately model the relation between a morph of e and the set of morphs in f to which it is connected. Model 4 is a more sophisticated step in this direction. Both model 3 and Model 4 are deficient. Model 5 is a generalization of Model 4 in this deficiency is removed at the expense of more increased complexity. For Models 3,4 and 5 the exact sum over alignments can not be computed efficiently. Instead, this sum can be approximated by restricting it to alignments of high probability.

Detailed Description Text - DETX (528):

In this section embodiments of the statistical translation model that assigns a conditional probability to the event that a sequence of lexical units in the source language is a translation of a sequence of lexical units in the target language will be described. Methods for estimating the parameters of these embodiments will be explained.

Detailed Description Text - DETX (532):

Random variables will be denoted by upper case letters, and the values of such variables will be denoted by the corresponding lower case letters. For random variables X and Y , the probability $\Pr(Y=y \text{ vertline } X=x)$ will be denoted by the shorthand $P(y \text{ vertline } x)$. Strings or vectors will be denoted by bold face letters, and their entries will be denoted by the corresponding non-bold letters.

Detailed Description Text - DETX (538):

Some embodiments of the translation model involve the idea of an alignment between a pair of strings. An alignment consists of a set of connections between words in the French string and words in the English string. An alignment is a random variable, A ; a generic values of this variable will be denoted by a . Alignments are shown graphically, as in FIG. 33, by writing the English string above the French string and drawing lines from some of the words in the English string to some of the words in the French string. These lines are called connections. The alignment in FIG. 33 has seven connections, (the, Le), (program, programme), and so on. In the description that follows, such alignment will be denoted as (Le programme a ete mis en application. vertline. And the(1) program(2) has(3)been(4) implemented(5,6,7)). The list of numbers following an English word shows the positions in the French string of the words with which it is aligned. Because And is not aligned with any French words here, there is no list of numbers after it. Every alignment is considered to be correct with some probability. Thus (Le programme a ete mis en application. vertline. And(1,2,3,4,5,6,7) the program has been implemented) is perfectly acceptable. Of course, this is much less probable than the alignment shown in FIG. 33.

Detailed Description Text - DETX (540):

The set of English words connected to a French word will be called the notion that generates it. An alignment

resolves the English string into a set of possibly overlapping notions that is called a notional scheme. The previous example contains the three notions. The, poor, and don't have any money. Formally, a notion is a subset of the positions in the English string together with the words occupying those positions. To avoid confusion in describing such cases, a subscript will be affixed to each word showing its position. The alignment in FIG. 34 includes the notions the.sub.4 and of.sub.6 the.sub.7, but not the notions of.sub.6 the.sub.4 or the.sub.7. In (J'applaudis a la decision.vertline.l(1) applaud(2) the(4) decision(5)), a is generated by the empty notion. Although the empty notion has no position and so never requires a clarifying subscript, it will be placed at the beginning of the English string, in position zero, and denoted by e.sub.0. At times, therefore, the previous alignment will be denoted as (J'applaudis a la decision.vertline. e.sub.0 (3) l(1) applaud(2) the(4) decision(5)).

Detailed Description Text - DETX (541):

The set of all alignments of (f.vertline.e) will be written (e,f). Since e has length l and f has length m, there are lm different connections that can be drawn between them. Since an alignment is determined by the connections that it contains, and since a subset of possible connections can be chosen in $2^{\text{sup.}lm}$ ways, there are $2^{\text{sup.}lm}$ alignments in (e,f).

Detailed Description Text - DETX (543):

The probability of a French string f and an alignment a given an English string e can be written ##EQU28## Here, m is the length of f and a.sub.1.sup.m is determined by a.

Detailed Description Text - DETX (545):

The conditional probabilities on the right-hand side of Equation (47) cannot all be taken as independent parameters because there are too many of them. In Model 1, the probabilities P(m.vertline.e) are taken to be independent of e and m; that P(a.sub.j.vertline.a.sub.1.sup.j-1, f.sub.1.sup.j-1, m, e), depends only on l, the length of the English string, and therefore must be $(l+1)^{\text{sup.}-1}$; and that P(f.sub.j.vertline.a.sub.1.sup.3,f.sub.1.sup.j-1, m, e) depends only on f.sub.j and e.sub.aj. The parameters, then, are .epsilon.tbd.P(m.vertline.e), and t(f.sub.j.vertline.e.sub.aj).tbd.P(f.sub.j.vertline.a.sub.1.sup.j, f.sub.1.sup.j-1, m,e), which will be called the translation probability of f.sub.j given e.sub.aj. The parameter .epsilon. is fixed at some small number. The distribution of M is unnormalized but this is a minor technical issue of no significance to the computations. In particular, M can be restricted to some finite range. As long as this range encompasses everything that actually occurs in training data, no problems arise.

Detailed Description Text - DETX (546):

A method for estimating the translation probabilities for Model 1 will now be described. The joint likelihood of a French sentence and an alignment is ##EQU29## The alignment is determined by specifying the values of a.sub.j for j from 1 to m, each of which can take any value from 0 to l. Therefore, ##EQU30##

Detailed Description Text - DETX (547):

The first goal of the training process is to find the values of the translation probabilities that maximize P(f.vertline.e) subject to the constraints that for each e, ##EQU31## An iterative method for doing so will be described.

Detailed Description Text - DETX (548):

The method is motivated by the following consideration. Following standard practice for constrained maximization, a necessary relationship for the parameter values at a local maximum can be found by introducing Lagrange multipliers, .lambda.sub.e, and seeking an unconstrained maximum of the auxiliary function ##EQU32## At a local maximum, all of the partial derivatives of h with respect to the components of t and .lambda. are zero. That the partial derivatives with respect to the components of .lambda. be zero is simply a restatement of the constraints on the translation probabilities. The partial derivative of h with respect to t (f.vertline.e) is ##EQU33## where .delta. is the Kronecker delta function, equal to one when both of its arguments are the same and equal to zero otherwise. This will be zero provided that ##EQU34## Superficially Equation (53)

looks like a solution to the maximization problem, but it is not because the translation probabilities appear on both sides of the equal sign. Nonetheless, it suggests an iterative procedure for finding a solution:

Detailed Description Text - DETX (549):

1. Begin with initial guess for the translation probabilities;

Detailed Description Text - DETX (553):

(Here and elsewhere, the Lagrange multipliers simply serve as a reminder that the translation probabilities must be normalized so that they satisfy Equation (50).) This process, when applied repeatedly is called the EM process. That it converges to a stationary point of h in situations like this, as demonstrated in the previous section, was first shown by L. E. Baum in an article entitled, An Inequality and Associated Maximization Technique in Statistical Estimation of Probabilistic Functions of a Markov Process, appearing the journal Inequalities, Vol. 3, in 1972.

Detailed Description Text - DETX (555):

In practice, the training data consists of a set of translations, $(f^{(1)} \cdot e^{(1)}), (f^{(2)} \cdot e^{(2)}), \dots, (f^{(S)} \cdot e^{(S)})$, so this equation becomes **##EQU37##** Here, λ_e serves only as a reminder that the translation probabilities must be normalized.

Detailed Description Text - DETX (556):

Usually, it is not feasible to evaluate the expectation in Equation (55) exactly. Even if multiword notions are excluded, there are still $(l+1) \cdot m$ alignments possible for $(f \cdot e)$. Model 1, however, is special because by recasting Equation (49), it is possible to obtain an expression that can be evaluated efficiently. The right-hand side of Equation (49) is a sum of terms each of which is monomial in the translation probabilities. Each monomial contains m translation probabilities, one for each of the words in f . Different monomials correspond to different ways of connecting words in f to notions in e with every way appearing exactly once. By direct evaluation, then **##EQU38##** Therefore, the sums in Equation (49) can be interchanged with the product to obtain **##EQU39##** Using this expression, it follows that **##STR1##** Thus, the number of operations necessary to calculate a count is proportional to lm rather than to $(l+1) \cdot m$ as Equation (55) might suggest.

Detailed Description Text - DETX (557):

The details of the initial guesses for $t(f \cdot e)$ are unimportant because $P(f \cdot e)$ has a unique local maximum for Model 1, as is shown in Section 10. In practice, the initial probabilities $t(f \cdot e)$ are chosen to be equal, but any other choice that avoids zeros would lead to the same final solution.

Detailed Description Text - DETX (559):

Model 1, takes no cognizance of where words appear in either string. The first word in the French string is just as likely to be connected to a word at the end of the English string as to one at the beginning. In contrast, Model 2 makes the same assumptions as in Model 1 except that it assumes that $P(a_j \cdot a_1^{j-1}, f_{1^{j-1}}^{j-1}, m, e)$ depends on j , a_j , and m , as well as on l . This is done using a set of alignment probabilities.

Detailed Description Text - DETX (562):

It is easily verified that Equations (54), (56), and (57) carry over from Model 1 to Model 2 unchanged. In addition, an iterative update formulas for $a(i \cdot j, m, l)$, can be obtained by introducing a new count, $c(i \cdot j, m, l; e)$, the expected number of times that the word in position j of f is connected to the word in position i of e . Clearly, **##EQU43##** In analogy with Equations (56) and (57), for a single translation, **##EQU44##** and, for a set of translations, **##EQU45##** Notice that if $f^{(s)}$ does not have the length m or if $e^{(s)}$ does not have length l , then the corresponding count is zero. As with the λ 's in earlier equations, the μ 's here serve to normalize the alignment probabilities.

Detailed Description Text - DETX (563):

Model 2 shares with Model 1 the important property that the sums in Equations (55) and (65) can be obtained efficiently. Equation (63) can be rewritten **##EQU46##** Using this form for $P(f_{\text{vertline}} e)$, it follows that **##EQU47##** Equation (69) has a double sum rather than the product of two single sums, as in Equation (60), because, in Equation (69), i and j are tied together through the alignment probabilities.

Detailed Description Text - DETX (564):

Model 1 is a special case of Model 2 in which $a(i_{\text{vertline}} j, m, l)$ is held fixed at $(l+1)_{\text{sup}}^{-1}$. Therefore, any set of parameters for Model 1 can be reinterpreted as a set of parameters for Model 2. Taking as initial estimates of the parameters for Model 2 the parameter values that result from training Model 1 is equivalent to computing the probabilities of all alignments using Model 1, but then collecting the counts appropriate to Model 2. The idea of computing the probabilities of the alignments using one model, but collecting the counts in a way appropriate to a second model is very general and can always be used to transfer a set of parameters from one model to another.

Detailed Description Text - DETX (566):

Models 1 and 2 make various approximations to the conditional probabilities that appear in Equation (47). Although Equation (47) is an exact statement, it is only one of many ways in which the joint likelihood of f and a can be written as a product of conditional probabilities. Each such product corresponds in a natural way to a generative process for developing f and a from e . In the process corresponding to Equation (47), a length for f is chosen first. Next, a position in e is selected and connected to the first position in f . Next, the identity of $f_{\text{sub}} 1$ is selected. Next, another position in e is selected and this is connected to the second word in f , and so on.

Detailed Description Text - DETX (575):

Model 3 is based on Equation (71). It assumes that, for i between 1 and l , $P(o_{\text{sub}} i_{\text{vertline}} o_{\text{sub}} 1_{\text{sup}} i-1, e)$ depends only on $o_{\text{sub}} i$ and $e_{\text{sub}} i$; that, for all i , $P(\tau_{\text{sub}} ik_{\text{vertline}} \tau_{\text{sub}} i1_{\text{sup}} k-1, \tau_{\text{sub}} 0_{\text{sup}} i-1, o_{\text{sub}} 0_{\text{sup}} l, e)$ depends only on $\tau_{\text{sub}} ik$ and $e_{\text{sub}} i$; and that, for i between 1 and l , $P(\pi_{\text{sub}} ik_{\text{vertline}} \pi_{\text{sub}} i1_{\text{sup}} k-1, \pi_{\text{sub}} 1_{\text{sup}} i-1, \tau_{\text{sub}} 0_{\text{sup}} l, o_{\text{sub}} 0_{\text{sup}} l, e)$ depends only on $\pi_{\text{sub}} ik$, i , m , and l . The parameters of Model 3 are thus a set of fertility probabilities, $n(o_{\text{vertline}} e_{\text{sub}} i)_{\text{td}} P(o_{\text{vertline}} o_{\text{sub}} 1_{\text{sup}} i-1, e)$; a set of translation probabilities, $t(f_{\text{vertline}} e_{\text{sub}} i)_{\text{td}} Pr(T_{\text{sub}} ik = f_{\text{vertline}} \tau_{\text{sub}} i1_{\text{sup}} k-1, \tau_{\text{sub}} 0_{\text{sup}} i-1, o_{\text{sub}} 0_{\text{sup}} l, e)$; and a set of distortion probabilities, $d(j_{\text{vertline}} i, m, l)_{\text{td}} Pr(l_{\text{sub}} ik = j_{\text{vertline}} \pi_{\text{sub}} i1_{\text{sup}} k-1, \pi_{\text{sub}} 1_{\text{sup}} i-1, \tau_{\text{sub}} 0_{\text{sup}} l, o_{\text{sub}} 0_{\text{sup}} l, e)$.

Detailed Description Text - DETX (576):

The distortion and fertility probabilities for $e_{\text{sub}} 0$ are handled differently. The empty notion conventionally occupies position 0, but actually has no position. Its purpose is to account for those words in the French string that cannot readily be accounted for by other notions in the English string. Because these words are usually spread uniformly throughout the French string, and because they are placed only after all of the other words in the sentence have been placed, the probability $Pr(l_{\text{sub}} Ok+1 = j_{\text{vertline}} \pi_{\text{sub}} 01_{\text{sup}} k, \pi_{\text{sub}} 1_{\text{sup}} l, \tau_{\text{sub}} 0_{\text{sup}} l, o_{\text{sub}} 0_{\text{sup}} l, e)$ is set to 0 unless position j is vacant in which case it is set $(o_{\text{sub}} 0 - k)_{\text{sup}}^{-1}$. Therefore, the contribution of the distortion probabilities for all of the words in $\tau_{\text{sub}} 0$ is $1/o_{\text{sub}} 0 !$.

Detailed Description Text - DETX (577):

The value of $o_{\text{sub}} 0$ depends on the length of the French sentence since longer sentences have more of these extraneous words. In fact Model 3 assumes that **##EQU51##** for some pair of auxiliary parameters $p_{\text{sub}} 0$ and $p_{\text{sub}} 1$. The expression on the left-hand side of this equation depends on $o_{\text{sub}} 1_{\text{sup}} l$ only through the sum $o_{\text{sub}} 1 + \dots + o_{\text{sub}} l$ and defines a probability distribution over $o_{\text{sub}} 0$ whenever $p_{\text{sub}} 0$ and $p_{\text{sub}} 1$ are nonnegative and sum to 1. The probability $P(o_{\text{sub}} 0_{\text{vertline}} o_{\text{sub}} 1_{\text{sup}} l, e)$ can be interpreted as follows. Each of the words from $\tau_{\text{sub}} 1_{\text{sup}} l$ is imagined to require an extraneous word with probability $p_{\text{sub}} 1$; this word is required to be connected to the empty notion. The probability that exactly $o_{\text{sub}} 0$ of the words from $\tau_{\text{sub}} 1_{\text{sup}} l$ will require an extraneous word is just the expression given in Equation (73).

Detailed Description Text - DETX (578):

As in Models 1 and 2, an alignment of (f.vertline.e) in Model 3 is determined by specifying a.sub.j for each position in the French string. The fertilities, o.sub.0 through o.sub.l, are functions of the a.sub.j's. Thus, o.sub.i is equal to the number of j's for which a.sub.j equals i. Therefore, ##EQU52## with .SIGMA..sub.f t(f.vertline.e)=1, .SIGMA..sub.j d(j.vertline.i, m, l)=1, .SIGMA..sub.o n(o.vertline.e)=1, and p.sub.0 + p.sub.1 =1. According to the assumptions of Model 3, each of the pairs (.tau.,.pi.) in (f,a) makes an identical contribution to the sum in Equation (72). The factorials in Equation (74) come from carrying out this sum explicitly. There is no factorial for the empty notion because it is exactly cancelled by the contribution from the distortion probabilities.

Detailed Description Text - DETX (582):

To define the subset, of the elements of (f.vertline.e) over which the sums are evaluated a little more notation is required. Two alignments, a and a' will be said to differ by a move if there is exactly one value of j for which a.sub.j .noteq.a.sub.j'. Alignments will be said to differ by a swap if a.sub.j =a.sub.j' except at two values, j.sub.1 and j.sub.2, for which a.sub.j1 =a.sub.j2' and a.sub.j2 =a.sub.j1'. The two alignments will be said to be neighbors if they are identical or differ by a move or by a swap. The set of all neighbors of a will be denoted by (a).

Detailed Description Text - DETX (583):

Let b(a) be that neighbor of a for which the likelihood is greatest. Suppose that ij is pegged for a. Among the neighbors of a for which ij is also pegged, let b.sub.i.rarw.j (a) be that for which the likelihood is greatest. The sequence of alignments a, b(a), b.sup.2 (a).quadbond.b(b(a)), . . . , converges in a finite number of steps to an alignment that will be denoted as b.sup.infin. (a). Similarly, if ij is pegged for a, the sequence of alignments a, b.sub.i.rarw.j (a), b.sub.i.rarw.j.sup.2 (a), . . . , converges in a finite number of steps to an alignment that will be denoted as b.sub.i.rarw.j.sup.infin. (a). The simple form of the distortion probabilities in Model 3 makes it easy to find b(a) and b.sub.i.rarw.j (a). If a' is a neighbor of a obtained from it by the move of j from i to i', and if neither i nor i' is 0, then ##EQU56## Notice that o.sub.i, is the fertility of the word in position i' for alignment a. The fertility of this word in alignment a' is o.sub.i' +1. Similar equations can be easily derived with either i or i' is zero, or when a and a' differ by a swap.

Detailed Description Text - DETX (586):

In one iteration of the EM process for Model 3, the counts in Equations (76) through (80), are computed by summing only over elements of . These counts are then used in Equations (81) through (84) to obtain a new set of parameters. If the error made by including only some of the elements of (e,f) is not too great, this iteration will lead to values of the parameters for which the likelihood of the training data is at least as large as for the first set of parameters.

Detailed Description Text - DETX (587):

The initial estimates of the parameters of Model 2 are adapted from the final iteration of the EM process for Model 2. That is, the counts in Equations (76) through (80) are computed using Model 2 to evaluate P (a.vertline.e,f). The simple for of Model 2 again makes exact calculation feasible. The Equations (69) and (70) are readily adapted to compute counts for the translation and distortion probabilities; efficient calculation of the fertility counts is more involved. A discussion of how this is done is given in Section 10.

Detailed Description Text - DETX (589):

A problem with the parameterization of the distortion probabilities in Model 3 is this: whereas the sum over all pairs .tau.,.pi. of the expression on the right-hand side of Equation (71) is unity, if Pr(lI.sub.ik =j.vertline..pi..sub.i1.sup.k-1,.pi..sub.1.sup.i-1,.tau..sub.0.sup.l, o.sub.0.sup.l,e) depends only on j, i, m, and l for i>0. Because the distortion probabilities for assigning positions to later words do not depend on the positions assigned to earlier words, Model 3 wastes some of its probability on what will be called generalized strings, i.e., strings that have some positions with several words and others with none. When a model has this property of not concentrating all of its probability on events of interest, it will be said to be deficient. Deficiency is the price for the simplicity that permits Equation (85).

Detailed Description Text - DETX (590):

Deficiency poses no serious problem here. Although Models 1 and 2 are not technically deficient, they are surely spiritually deficient. Each assigns the same probability to the alignments (Je n'ai pas de stylo.vertline.l(1) do not(2,4) have(3) a(5) pen(6)) and (Je pas ai ne de stylo.vertline.l(1) do not(2,4) have(3) a(5) pen(6)), and, therefore, essentially the same probability to the translations (Je n'ai pas de stylo.vertline.l do not have a pen) and (Je pas ai ne de style.vertline.l do not have a pen). In each case, not produces two words, ne and pas, and in each case, one of these words ends up in the second position of the French string and the other in the fourth position. The first translation should be much more probable than the second, but this defect is of little concern because while the system may be required to translate the first string someday, it will almost surely not be required to translate the second. The translation models are not used to predict French given English but rather as a component of a system designed to predict English given French. They need only be accurate to within a constant factor over well-formed strings of French words.

Detailed Description Text - DETX (592):

Often the words in an English string constitute phrases that are translated as units into French. Sometimes, a translated phrase may appear at a spot in the French string different from that at which the corresponding English phrase appears in the English string. The distortion probabilities of Model 3 do not account well for this tendency of phrases to move around as units. Movement of a long phrase will be much less likely than movement of a short phrase because each word must be moved independently. In Model 4, the treatment of $\Pr(l_{\text{sub}.ik} = j_{\text{vertline.}} \cdot \pi_{\text{sub}.i1} \cdot \sup.k-1, \pi_{\text{sub}.1} \cdot \sup.i-1, \tau_{\text{sub}.0} \cdot \sup.l, o_{\text{sub}.0} \cdot \sup.l, e)$ is modified so as to alleviate this problem. Words that are connected to the empty notion do not usually form phrases and so Model 4 continues to assume that these words are spread uniformly throughout the French string.

Detailed Description Text - DETX (594):

In Model 4, the probabilities $d(j_{\text{vertline.}}i, m, l)$ are replaced by two sets of parameters: one for placing the head of each notion, and one for placing any remaining words. For $[i] > 0$, Model 4 requires that the head for notion i be $\tau_{\text{sub}.[i]1}$. It assumes that

Detailed Description Text - DETX (595):

Here, d and e are functions of the English and French word that take on a small number of different values as their arguments range over their respective vocabularies. In the Section entitled Classes, a process is described for dividing a vocabulary into classes so as to preserve mutual information between adjacent classes in running text. The classes d and e are constructed as functions with fifty distinct values by dividing the English and French vocabularies each into fifty classes according to this method. The probability is assumed to depend on the previous notion and on the identity of the French word being placed so as to account for such facts as the appearance of adjectives before nouns in English but after them in French. The displacement for the head of notion i is denoted by $j_{\text{sub}.i-1}$. It may be either positive or negative. The probability of $d_{\text{sub}.1}(-1_{\text{vertline.}}(e), (f))$ is expected to be larger than $d_{\text{sub}.1}(+1_{\text{vertline.}}(e), (f))$ when e is an adjective and f is a noun. Indeed, this is borne out in the trained distortion probabilities for Model 4, where $d_{\text{sub}.1}(-1_{\text{vertline.}}(\text{government's}), (\text{developpement}))$ is 0.9269, while $d_{\text{sub}.1}(+1_{\text{vertline.}}(\text{government's}), (\text{developpement}))$ is 0.0022.

Detailed Description Text - DETX (599):

The count and reestimation formulae for Model 4 are similar to those for the previous Models and will not be given explicitly here. The general formulae in Section 10 are helpful in deriving these formulae. Once again, the several counts for a translation are expectations of various quantities over the possible alignments with the probability of each alignment computed from an earlier estimate of the parameters. As with Model 3, these expectations are computed by sampling some small set, \mathcal{A} , of alignments. As described above, the simple form that for the distortion probabilities in Model 3, makes it possible to find $b_{\text{sup}.infin.}(a)$ rapidly for any a . The analogue of Equation (85) for Model 4 is complicated by the fact that when a French word is moved from notion to notion, the centers of two notions change, and the contribution of several words is affected. It is nonetheless possible to evaluate the adjusted likelihood incrementally, although it is substantially more time consuming.

Detailed Description Text - DETX (607):

Again, the final factor enforces the constraint that τ_{ijk} land in a vacant position, and, again, it is assumed that the probability depends on $f_{sub,j}$ only through its class.

Detailed Description Text - DETX (611):

A large collection of training data is used to estimate the parameters of the five models described above. In one embodiment of these models, training data is obtained using the method described in detail in the paper, Aligning Sentences in Parallel Corpora, by P. F. Brown, J. C. Lai, and R. L. Mercer, appearing in the Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics, June 1991. This paper is incorporated by reference herein. This method is applied to a large number of translations from several years of the proceedings of the Canadian parliament. From these translations, a training data set is chosen comprising those pairs for which both the English sentence and the French sentence are thirty words or less in length. This is a collection of 1,778,620 translations. In an effort to eliminate some of the typographical errors that abound in the text, a English vocabulary is chosen consisting of all of those words that appear at least twice in English sentences in the data, and as a French vocabulary is chosen consisting of all those words that appear at least twice in French sentences in the data. All other words are replaced with a special unknown English word or unknown French word according as they appear in an English sentence or a French sentence. In this way an English vocabulary of 42,005 words and a French vocabulary of 58,016 words is obtained. Some typographic errors are quite frequency, for example, momento for memento, and so the vocabularies are not completely free of them. At the same time, some words are truly rare, and in some cases, legitimate words are omitted. Adding $e_{sub,0}$ to the English vocabulary brings it to 42,006 words.

Detailed Description Text - DETX (612):

Eleven iterations of the EM process are performed for this data. The process is initialized by setting each of the 2,437,020,096 translation probabilities, $t(f_{sub,j}, e_{sub,i})$, to $1/58016$. That is, each of the 58,016 words in the French vocabulary is assumed to be equally likely as a translation for each of the 42,006 words in the English vocabulary. For $t(f_{sub,j}, e_{sub,i})$ to be greater than zero at the maximum likelihood solution for one of the models, f and e must occur together in at least one of the translations in the training data. This is the case for only 25,427,016 pairs, or about one percent of all translation probabilities. On the average, then, each English word appears with about 605 French words.

Detailed Description Text - DETX (613):

Table 6 summarizes the training computation. At each iteration, the probabilities of the various alignments of each translation using one model are computed, and the counts using a second, possibly different model are accumulated. These are referred to in the table as the In model and the Out model, respectively. After each iteration, individual values are retrained for only those translation probabilities that surpass a threshold; the remainder are set to the small value (10^{-12}). This value is so small that it does not affect the normalization conditions, but is large enough that translation probabilities can be resurrected during later

Detailed Description Text - DETX (614):

iterations. As is apparent from columns 4 and 5, even though the threshold is lowered as iterations progress, fewer and fewer probabilities survive. By the final iteration, only 1,620,287 probabilities survive, an average of about thirty-nine French words for each English word.

Detailed Description Text - DETX (615):

As has been described, when the In model is neither Model 1 nor Model 2, the counts are computed by summing over only some of the possible alignments. Many of these alignments have a probability much smaller than that of the Viterbi alignment. The column headed Alignments in Table 6 shows the average number of alignments for which the probability is within a factor of 25 of the probability of the Viterbi alignment in each iteration. As this number drops, the model concentrates more and more probability onto fewer and fewer alignments so that the Viterbi alignment becomes ever more dominant.

Detailed Description Text - DETX (616):

The last column in the table shows the perplexity of the French text given the English text for the In model of the iteration. The likelihood of the training data is expected to increase with each iteration. This likelihood can be thought of as arising from a product of factors, one for each French word in the training data. There are 28,850,104 French words in the training data so the 28,850,104^{sup.th} root of the likelihood is the average factor by which the likelihood is reduced for each additional French word. The reciprocal of this root is the perplexity shown in the table. As the likelihood increases, the perplexity decreases. A steady decrease in perplexity is observed as the iterations progress except when a switch from Model 2 as the In model to Model 3 is made. This sudden jump is not because Model 3 is a poorer model than Model 2, but because Model 3 is deficient; the great majority of its probability is squandered on objects that are not strings of French words. As has been explained, deficiency is not a problem. In the description of Model 1, the $P(m.\text{vertline}.e)$ was left unspecified. In quoting perplexities for Models 1 and 2, it is assumed that the length of the French string is Poisson with a mean that is a linear function of the length of the English string. Specifically, it is assumed that $\Pr(M=m.\text{vertline}.e) = (\lambda^m / m!) \cdot e^{-\lambda}$, with λ equal to 1.09.

Detailed Description Text - DETX (618):

Tables 8 through 17 show the translation probabilities and fertilities after the final iteration of training for a number of English words. All and only those probabilities that are greater than 0.01 are shown. Some words, like nodding, in Table 8, do not slip gracefully into French. Thus, there are translations like (Il fait signe que oui.vertline.He is nodding), (Il fait un signe de la tete.vertline.He is nodding), (Il fait un signe de tete affirmatif.vertline.He is nodding), or (Il hoche la tete affirmativement.vertline.He is nodding). As a result, nodding frequently has a large fertility and spreads its translation probability over a variety of words. In French, what is worth saying is worth saying in many different ways. This is also seen with words like should, in Table 9, which rarely has a fertility greater than one but still produces many different words, among them devrait, devraient, devrions, doit, doivent, devons, and devrais. These are (just a fraction of the many) forms of the French verb devoir. Adjectives fare a little better: national, in Table 10, almost never produces more than one word and confines itself to one of nationale, national, nationaux, and nationales, respectively the feminine, the masculine, the masculine plural, and the feminine plural of the corresponding French adjective. It is clear that the models would benefit from some kind of morphological processing to rein in the lexical exuberance of French.

Detailed Description Text - DETX (620):

together with an article. Thus, farmers typically has a fertility 2 and usually produces either agriculteurs or les. Additional examples are included in Tables 13 through 17, which show the translation and fertility probabilities for external, answer, oil, farmer, and not.

Detailed Description Text - DETX (621):

FIGS. 37, 38, and 39 show automatically derived alignments for three translations. In the terminology used above, each alignment is $b.\text{sup}.\text{infin}.\text{V}(f.\text{vertline}.e;2)$. It should be understood that these alignments have been found by a process that involves no explicit knowledge of either French or English. Every fact adduced to support them has been discovered automatically from the 1,778,620 translations that constitute the training data. This data, in turn, is the product of a process the sole linguistic input of which is a set of rules explaining how to find sentence boundaries in the two languages. It may justifiably be argued, therefore, that these alignments are inherent in the Canadian Hansard data itself.

Detailed Description Text - DETX (624):

The final example, in FIG. 39, has several features that bear comment. The second word, Speaker, is connected to the sequence l'Orateur. Like farmers above, it has trained to produce both the word that one naturally thinks of as its translation and the article in front of it. In the Hansard data, Speaker always has fertility 2 and produces equally often l'Orateur and le president. Later in the sentence, starred is connected to the phrase marquées de un asterisque. From an initial situation in which each French word is equally probable as a translation of starred, the system has arrived, through training, at a situation where it is able to connect to just the right string of four words. Near the end of the sentence, give is connected to donnerai, the first person singular future of donner, which means to give. It might be better if both will and give were connected to donnerai, but by limiting notions to no more than one word, Model 5 precludes this possibility. Finally, the last twelve words of the English sentence, I now have the answer and will give it to the House, clearly correspond to the last seven words of the French sentence, je donnerai la reponse a la Chambre, but, literally, the French is I will give the answer to

the House. There is nothing about not, have, and, or it, and each of these words has fertility 0. Translations that are as far as this from the literal are rather more the rule than the exception in the training data.

Detailed Description Text - DETX (629):

It has been argued above that neither spiritual nor actual deficiency poses a serious problem, but this is not entirely true. Let $w(e)$ be the sum of $P(f.\text{verline}.e)$ over well-formed French strings and let $i(e)$ be the sum over ill-formed French strings. In a deficient model, $w(e)+i(e)<1$. In this case, the remainder of the probability is concentrated on the event failure and so $w(e)+i(e)+P(\text{failure}.\text{verline}.e)=1$. Clearly, a model is deficient precisely when $P(\text{failure}.\text{verline}.e)>0$. if $P(\text{failure}.\text{verline}.e)=0$, but $i(e)>0$, then the model is spiritually deficient. If $w(e)$ were independent of e , neither form of deficiency would pose a problem, but because Models 1-5 have no long-term constraints, $w(e)$ decreases exponentially with l . When computing alignments, even this creates no problem because e and f are known. However, for a given f , if the goal is to discover the most probable e , then the product $P(e)P(f.\text{verline}.e)$ is too small for long English strings as compared with short ones. As a result, short English strings are improperly favored over longer English strings. This tendency is counteracted in part by the following modification:

Detailed Description Text - DETX (633):

Models 1 through 5 all assign non-zero probability only to alignments with notions containing no more than one word each. Except in Models 4 and 5, the concept of a notion plays little role in the development. Even in these models, notions are determined implicitly by the fertilities of the words in the alignment: words for which the fertility is greater than zero make up one-word notions; those for which it is zero do not. It is not hard to give a method for extending the generative process upon which Models 3, 4, and 5 are based to encompass multi-word notions. This method comprises the following enhancements:

Detailed Description Text - DETX (645):

$\text{.epsilon}.(m.\text{verline}.l)$ string length probabilities

Detailed Description Text - DETX (646):

$t(f.\text{verline}.e)$ translation probabilities

Detailed Description Text - DETX (653):

1. Choose a length m for f according to the probability distribution $\text{.epsilon}.(m.\text{verline}.l)$.

Detailed Description Text - DETX (664):

$\text{.epsilon}.(m.\text{verline}.l)$ string length probabilities

Detailed Description Text - DETX (665):

$t(f.\text{verline}.e)$ translation probabilities

Detailed Description Text - DETX (666):

$a(i.\text{verline}.j, l, m)$ alignment probabilities

Detailed Description Text - DETX (670):

This model is not deficient. Model 1 is the special case of this model in which the alignment probabilities are uniform: $a(i.\text{verline}.j, l, m)=(l+1)^{-1}$ for all i .

Detailed Description Text - DETX (744):

N.B. In the previous section, simplified embodiment of this model in which the probabilities $d_{\text{sub.LR}}$ do not appear is described.

Detailed Description Text - DETX (781):

4201. For the word being considered, finding a good question for each of a plurality of informant sites. These informant sites are obtained from a table 4207 stored in memory of possible informant sites. Possible sites include but are not limited to, the nouns to the right and left, the verbs to the right and left, the words to the right and left, the words two positions to the right or left, etc. A method of finding a good question is described below. This method makes use of a table 4205 stored in memory probabilities derived from Viterbi alignments. These probabilities are also discussed below.

Detailed Description Text - DETX (795):

The purpose of source and target transduction is to facilitate the task of the statistical translation. This will be accomplished if the probability distribution $\Pr(f', e')$ is easier to model than the original distribution $\Pr(f, e)$. In practice this means that e' and f' should encode global linguistic facts about e and f in a local form.

Detailed Description Text - DETX (797):

The cross entropy measures the average uncertainty that the model has about the target language translation e of a source language sequence f . Here $P(e \cdot \text{verline} \cdot f)$ is the probability according to the model that e is a translation of f and the sum runs over a collection of all S pairs of sentences in a large corpus comprised of pairs of sentences with each pair consisting of a source and target sentence which are translations of one another (See Sections 8-10).

Detailed Description Text - DETX (803):

The probability models. A translation model such as one of the models in Sections 8-10 is used for both $P'(F' \cdot \text{verline} \cdot E')$ and for $P(F \cdot \text{verline} \cdot E)$. A trigram language model such as that discussed in Section 6 is used for both $P(E)$ and $P'(E')$.

Detailed Description Text - DETX (805):

The probability $P(f \cdot \text{verline} \cdot e)$ computed by the translation model requires a sum over alignments as discussed in detail in Sections 8-10. This sum is often too expensive to compute directly since the number of alignments increases exponentially with sentence length. In the mathematical considerations of this Section, this sum will be approximated by the single term corresponding to the alignment, $(f \cdot \text{verline} \cdot e)$, with greatest probability. This is the Viterbi approximation already discussed in Sections 8-10 and $(f \cdot \text{verline} \cdot e)$ is the Viterbi alignment.

Detailed Description Text - DETX (806):

Let $c(f \cdot \text{verline} \cdot e)$ be the expected number of times that e is aligned with f in the Viterbi alignment of a pair of sentences drawn at random from a large corpus of training data. Let $(c \cdot \text{phi} \cdot \text{verline} \cdot e)$ be the expected number of times that e is aligned with phi words. Then ##EQU91## where $c(f \cdot \text{verline} \cdot e;)$ is the number of times that e is aligned with f in the alignment A , and $c(\text{phi} \cdot \text{verline} \cdot e;)$ is the number of times that e generates phi target words in A . The counts above are also expressible as averages with respect to the model: ##EQU92##

Detailed Description Text - DETX (807):

Probability distributions $p(e, f)$ and $p(\text{phi}, e)$ are obtained by normalizing the counts $c(f \cdot \text{verline} \cdot e)$ and $(c \cdot \text{phi} \cdot \text{verline} \cdot e)$: ##EQU93##

Detailed Description Text - DETX (808):

(These are the probabilities that are stored in a table of probabilities 4205.) The conditional distributions $p(f.\text{vertline}.e)$ and $p(\phi.\text{vertline}.e)$ are the Viterbi approximation estimates for the parameters of the model. The marginals satisfy $\sum_e u(e)$ where $u(e)$ and $u(f)$ are the unigram distributions of e and f and $\phi(e) = \sum_f p(\phi.\text{vertline}.e)$. ϕ is the average number of source words aligned with e . These formulae reflect the fact that in any alignment each source word is aligned with exactly one target word.

Detailed Description Text - DETX (812):

where m is the average length of the source sentences in the training data, and $H(F.\text{vertline}.E)$ and $H(\phi.\text{vertline}.E)$ are the conditional entropies for the probability distributions $p(f,e)$ and $p(\phi,e)$: $\sum_e H(\phi.\text{vertline}.e)$

Detailed Description Text - DETX (815):

where $H(E)$ is the cross entropy of $P(E)$ and $I(F,E)$ is the mutual information between f and e for the probability distribution $p(f,e)$.

Detailed Description Text - DETX (817):

Next consider $H(E'.\text{vertline}.F')$. Let $e.f\text{wdarw}.e'$ and $f.f\text{wdarw}.f'$ be sense labeling transformations of the type discussed above. Assume that these transformations preserve Viterbi alignments; that is, the words e and f are aligned in the Viterbi alignment for (f,e) , then their sensed versions e' and f' are aligned in the Viterbi alignment for (f',e') . It follows that the word translation probabilities obtained from the Viterbi alignments satisfy $p(f,e) = \sum_{f'} p(f',e) p(f,e')$, where the sums range over the sensed versions f' of f and the sensed versions e' of e .

Detailed Description Text - DETX (823):

For sensing source sentences, a question about an informant is a function c from the source vocabulary into the set of possible senses. If the informant of f is x , then f is assigned the sense $c(x)$. The function $c(x)$ is chosen to minimize the cross entropy $H(E.\text{vertline}.F')$. From formula (175), this is equivalent to maximizing the conditional mutual information $I(F'.\text{vertline}.f)$ between E and F' $\sum_x p(f,e,x)$ where $p(f,e,x)$ is the probability distribution obtained by counting the number of times in the Viterbi alignments that e is aligned with f and the value of the informant of f is x , $\sum_x p(f,e,x)$

Detailed Description Text - DETX (825):

The method is based on the fact that, up to a constant independent of c , the mutual information $I(f'.\text{vertline}.f)$ can be expressed as an infimum over conditional probability distributions $q(E.\text{vertline}.c)$, $\sum_x p(f,e,x)$

Detailed Description Text - DETX (836):

The methods of sense-labeling discussed above ask a single question about a single word of context. In other embodiments of the sense labeler, this question is the first question in a decision tree. In still other embodiments, rather than using a single informant site to determine the sense of a word, questions from several different informant sites are combined to determine the sense of a word. In one embodiment, this is done by assuming that the probability of an informant word $x.\text{sub}.i$ at informant site i , given a target word e , is independent of an informant word $x.\text{sub}.j$ at a different informant site j given the target word e . Also, in other embodiments, the intermediate source and target structure representations are more sophisticated than word sequences, including, but not limited to, sequences of lexical morphemes, case frame sequences, and parse tree structures.

Detailed Description Text - DETX (839):

A number of researchers have developed methods that align sentences according to the words that they contain. (See for example, Deriving translation data from bilingual text, by R. Catizone, G. Russel, and S. Warwick, appearing in Proceedings of the First International Acquisition Workshop, Detroit, Mich. 1989; and "Making Connections", by M. Kay, appearing in ACH/ALLC '91, Tempe, Ariz. 1991.) Unfortunately, these methods are necessarily slow and, despite the potential for high accuracy, may be unsuitable for very large collections of

text.

Detailed Description Text - DETX (856):

In the Hansard example, the English corpus contains 85,016,286 tokens in 3,510,744 sentences, and the French corpus contains 97,857,452 tokens in 3,690,425 sentences. The average English sentence has 24.2 tokens, while the average French sentence is about 9.5% longer with 26.5 tokens. The left-hand side of FIG. 48 shows the raw data for a portion of the English corpus, and the right-hand side shows the same portion after it was cleaned, tokenized, and divided into sentences. The sentence numbers do not advance regularly because the sample has been edited in order to display a variety of phenomena.

Detailed Description Text - DETX (866):

Given these costs, the standard technique of dynamic programming is used to find the alignment between the major anchors with the least total cost. Dynamic programming is described by R. Bellman in the book titled Dynamic Programming, published by Princeton University Press, Princeton, N.J. in 1957. In theory, the time and space required to find this alignment grow as the product of the lengths of the two sequences to be aligned. In practice, however, by using thresholds and the partial traceback technique described by Brown, Spohrer, Hochschild, and Baker in their paper, Partial Traceback and Dynamic Programming, published in the Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, in Paris, France in 1982, the time required can be made linear in the length of the sequences, and the space can be made constant. Even so, the computational demand is severe. In the Hansard example, the two corpora were out of alignment in places by as many as 90,000 sentences owing to mislabelled or missing files.

Detailed Description Text - DETX (872):

The generation of beads is modelled by the two-state Markov model shown in FIG. 50. The allowed beads are shown in FIG. 51. A single sentence in one corpus is assumed to line up with zero, one, or two sentences in the other corpus. The probabilities of the different cases are assumed to satisfy $\Pr(e)=\Pr(f)$, $\Pr(ef)=\Pr(eef)$, and $\Pr(.sub.e)=\Pr(.sub.f)$.

Detailed Description Text - DETX (873):

The generation of sentence lengths given beads is modeled as follows. The probability of an English sentence of length $l.sub.e$ given an e-bead is assumed to be the same as the probability of an English sentence of length $l.sub.e$ in the text as a whole. This probability is denoted by $\Pr(l.sub.e)$. Similarly, the probability of a French sentence of length $l.sub.f$ given an f-bead is assumed to equal $\Pr(l.sub.f)$. For an ef-bead, the probability of an English sentence of length $l.sub.e$ is assumed to equal $\Pr(l.sub.e)$ and the log of the ratio of length of the French sentence to the length of the English sentence is assumed to be normally distributed with mean μ and variance σ^2 . Thus, if $r=\log(l.sub.f / l.sub.e)$, then

Detailed Description Text - DETX (874):

with α chosen so that the sum of $\Pr(l.sub.f | l.sub.e)$ over positive values of $l.sub.f$ is equal to unity. For an eef-bead, the English sentence lengths are assumed to be independent with equal marginals $\Pr(l.sub.e)$, and the log of the ratio of the length of the French sentence to the sum of the lengths of the English sentences is assumed to be normally distributed with the same mean and variance as for an ef-bead. Finally, for an eef-bead, the probability of an English length $l.sub.e$ is assumed to equal $\Pr(l.sub.e)$ and the log of the ratio of the sum of the lengths of the French sentences to the length of the English sentence is assumed to be normally distributed as before. Then, given the sum of the lengths of the French sentences, the probability of a particular pair of lengths, $l.sub.f1$ and $l.sub.f2$, is assumed to be proportional to $\Pr(l.sub.f1)\Pr(l.sub.f2)$.

Detailed Description Text - DETX (875):

Together, the model for sequences of beads and the model for sentence lengths given beads define a hidden Markov model for the generation of aligned pairs of sentence lengths. Markov Models are described by L. Baum in the article "An Inequality and associated maximization technique in statistical estimation of probabilistic

functions of a Markov process", appearing in Inequalities in 1972.

Detailed Description Text - DETX (876):

The distributions $\text{Pr}(\text{l.sub.e})$ and $\text{Pr}(\text{l.sub.f})$ are determined from the relative frequencies of various sentence lengths in the data. For reasonably small lengths, the relative frequency is a reliable estimate of the corresponding probability. For longer lengths, probabilities are determined by fitting the observed frequencies of longer sentences to the tail of a Poisson distribution. The values of the other parameters of the Markov model can be determined by from a large

Detailed Description Text - DETX (883):

By repeating this process many thousands of times, an expected error rate of about 0.9% was estimated for the actual frequency of anchor points in the Hansard data. By varying the parameters of the hidden Markov model, the effect of anchor points and paragraph markers on error rate can be explored. With paragraph markers but no anchor points, the expected error rate is 2.0%, with anchor points but no paragraph markers, the expected error rate is 2.3%, and with neither anchor points nor paragraph markers, the expected error rate is 3.2%. Thus, while anchor points and paragraph markers are important, alignment is still feasible without them. This is promising since it suggests that the method is applicable to corpora for which frequent anchor points are not available.

Detailed Description Text - DETX (894):

Word-by-word alignments obtained in this way offer a valuable resource for work in bilingual lexicography and machine translation. For example, a method of cross-lingual sense labeling, described in Section 11, and also in the aforementioned paper, "Word Sense Disambiguation using Statistical Methods", uses alignments obtained in this way as data for construction of a statistical sense-labelling module.

Detailed Description Text - DETX (900):

A set of partial hypotheses is initialized in step 5401. A partial hypothesis is comprised of a target structure and an alignment with some subset of the morphemes in the source structure to be translated. The initial set generated by step 5401 consists of a single partial hypothesis. The partial target structure for this partial hypothesis is just an empty sequence of morphemes. The alignment is the empty alignment in which no morphemes in the source structure to be translated are accounted for.

Detailed Description Text - DETX (901):

The system then enters a loop through steps 5402, 5403, and 5404, in which partial hypotheses are iteratively extended until a test for completion is satisfied in step 5403. At the beginning of this loop, in step 5402, the existing set of partial hypotheses is examined and a subset of these hypotheses is selected to be extended in the steps which comprise the remainder of the loop. In step 5402 the score for each partial hypothesis is compared to a threshold (the method used to compute these thresholds is described below). Those partial hypotheses with scores greater than threshold are then placed on a list of partial hypotheses to be extended in step 5404. Each partial hypothesis that is extended in step 5404 contains an alignment which accounts for a subset of the morphemes in the source sentence. The remainder of the morphemes must still be accounted for. Each extension of an hypothesis in step 5404 accounts for one additional morpheme. Typically, there are many tens or hundreds of extensions considered for each partial hypothesis to be extended. For each extension a new score is computed. This score contains a contribution from the language model as well as a contribution from the translation model. The language model score is a measure of the plausibility a priori of the target structure associated with the extension. The translation model score is a measure of the plausibility of the partial alignment associated with the extension. A partial hypothesis is considered to be a full hypothesis when it accounts for the entire source structure to be translated. A full hypothesis contains an alignment in which every morpheme in the source structure is aligned with a morpheme in the hypothesized target structure. The iterative process of extending partial hypotheses terminates when step 5402 produces an empty list of hypotheses to be extended. A test for this situation is made on step 5403.

Detailed Description Text - DETX (910):

Here, * is a symbol which denotes a sequence boundary, and the factor $l(\text{the.vertline.*,*})$ is the trigram language model parameter that serves as an estimate of the probability that the English morpheme the occurs at the beginning of a sentence. The factor $n(1.\text{vertline.the})$ is the translation model parameter that is an estimate of the probability that the English morpheme the has fertility 1, in other words, that the English morpheme the is aligned with only a single French morpheme. The factor $t(\text{la.vertline.the})$ is the translation model parameter that serves as an estimate of the lexical probability that the English morpheme the translates to the French morpheme la. Finally, the factor $d(1.\text{vertline.1})$ is the translation model parameter that serves as an estimate of the distortion probability that a French morpheme will be placed in position 1 of the French structure given that it is aligned with an English morpheme that is in position 1 of the English structure. In the second example in FIG. 57, the English morpheme mother is hypothesized as accounting for the French morpheme mere. The score for this partial hypothesis is

Detailed Description Text - DETX (911):

Here, the final $d(7.\text{vertline.1})$ serves as an estimate of the distortion probability that a French morpheme, such as mere, will be placed in the 7th position in a source sequence given that it is aligned with an English morpheme such as mother which is in the 1st position in an hypothesized target sequence.

Detailed Description Text - DETX (913):

Here, the factor $l(\text{girl.vertline.*},\text{the})$ is the language model parameter that serves as an estimate of the probability with which the English morpheme girl is the second morpheme in a source structure in which the first morpheme is the. The next factor of 2 is the combinatorial factor that is discussed in the section entitled Translation Models and Parameter Estimation. It is factored in, in this case, because the open English morpheme girl is to be aligned with at least two French morphemes. The factor $n(i.\text{vertline.girl})$ is the translation model parameter that serves as an estimate of the probability that the English morpheme girl will be aligned with exactly i French morphemes, and the sum of these parameters for i between 2 and 25 is an estimate of the probability that girl will be aligned with at least 2 morphemes. It is assumed that the probability that an English morpheme will be aligned with more than 25 French morphemes is 0. Note that in a preferred embodiment of the present invention, this sum can be prepared and stored in memory as a separate parameter. The factor $t(\text{fille.vertline.girl})$ is the translation model parameter that serves as an estimate of the lexical probability that one of the French morphemes aligned with the English morpheme girl will be the French morpheme fille. Finally, the factor $d(3.\text{vertline.2})$ is the translation model parameter that serves as an estimate of the distortion probability that a French morpheme will be placed in position 3 of the French structure given that it is aligned with an English morpheme which is in position 2 of the English structure. This extension score in Equation 184 is multiplied by the score in Equation 182 for the partial hypothesis which is being extended to yield a new score for the partial hypothesis in FIG. 56 of ##EQU106##

Detailed Description Text - DETX (916):

Here, the first quotient adjusts the fertility score for the partial hypothesis by dividing out the estimate of the probability that girl is aligned with at least two French morphemes and by multiplying in an estimate of the probability that girl is aligned with exactly two French morphemes. As in the other examples, the second and third factors are estimates of the lexical and distortion probabilities associated with this extension.

Detailed Description Text - DETX (921):

Here, the first two factors are the trigram language model estimates of the probabilities with which up follows girl to.sub.— wake, and with which her follows to.sub.— wake up, respectively. The third factor is the fertility parameter that serves as an estimate of the probability that up is aligned with no source morphemes. The fourth, fifth, and sixth factors are the appropriate fertility, lexical, and distortion parameters associated with the target morpheme her in this partial alignment.

Detailed Description Text - DETX (923):

The score for this extension differs from the score in Equation 188 in that the fertility parameter $n(1.\text{vertline.her})$ is replaced by the combinatorial factor 2 and the sum of fertility parameters which provides an estimate of the probability that her will be aligned with at least two source morphemes.

Detailed Description Text - DETX (924):

A remaining type of extension is where a partial hypothesis is extended by an additional connection which aligns a source morpheme with the null target morpheme. The score for this type of extension is similar to those described above. No language model score is factored in, and scores from the translation model are factored in, in accordance with the probabilities associated with the null word as described in the section entitled Translation Models and Parameter Estimation.

Detailed Description Text - DETX (926):

Throughout the hypothesis search process, partial hypotheses are maintained in a set of priority queues. In theory, there is a single priority queue for each subset of positions in the source structure. So, for example, for the source structure oui, oui, three positions, oui is in position 1; a comma is in position 2; and oui is in position 3, and there are therefore 2.sup.3 subsets of positions: [], [1], [2], [3], [1,2], [1,3], [2,3], and [1,2,3]. In practice these priority queues are initialized only on demand, and many less than the full number of queues possible are used in the hypothesis search. In a preferred embodiment, each partial hypothesis is comprised of a sequence of target morphemes, and these morphemes are aligned with a subset of source morphemes. Corresponding to that subset of source morphemes is a priority queue in which the partial hypothesis is stored. The partial hypotheses within a queue are prioritized according to the scores associated with those hypotheses. In certain preferred embodiments the priority queues are limited in size and only the 1000 hypothesis with the best scores are maintained.

Detailed Description Text - DETX (927):

The set of all subsets of a set of source structure positions can be arranged in a subset lattice. For example, the subset lattice for the set of all sets of the set [1,2,3] is shown in FIG. 64. In a subset lattice, a parent of a set S is any which contains one less element than S, and which is also a subset of S. In FIG. 64 arrows have been drawn from each set in the subset lattice to each of its parents. For example, the set [2] is a parent of the set [1,2].

Detailed Description Text - DETX (928):

A subset lattice defines a natural partial ordering on a set of sets. Since the priority queues used in hypothesis search are associated with subsets, a subset lattice also defines a natural partial ordering on the set of priority queues. Thus in FIG. 64, there are two parents of the priority queue associated with the subset of source structure positions [1,3]. These two parents are the priority queues associated with the set [1] and [3]. A priority queue Q.sub.1 is said to be an ancestor of another priority Q.sub.2 if 1) Q.sub.1 is not equal to Q.sub.2, and 2) Q.sub.1 is a subset of Q.sub.2. If Q.sub.1 is an ancestor of Q.sub.2, then Q.sub.2 is said to be to be a descendant of Q.sub.1.

Detailed Description Text - DETX (930):

A priority queue is said to be active if there are partial hypotheses stored in it. An active priority queue is said to be on the frontier if it has no active descendent. The cardinality of a priority queue is equal to the number of elements in the subset with which it is associated. So, for example, the cardinality of the priority queue which is associated with the set [2,3] is 2.

Detailed Description Text - DETX (931):

The process in step 5402 functions by assigning a threshold to every active priority queue and then places on the list of partial hypotheses to be extended every partial hypothesis on an active priority queue that has a score that is greater than the threshold for that priority queue. This is depicted in FIG. 66. First, in step 6601 the threshold for every active priority queue is initialized to infinity, in practice, some very large number. Second, in step 6602, thresholds are determined for every priority queue on the frontier.

Detailed Description Text - DETX (932):

The method by which these thresholds are computed is best described by first describing what the normalizer of

a priority queue is. Each priority queue on the frontier corresponds to a set of positions of source morphemes. At each position of these positions is a particular source morpheme. Associated with each morpheme is a number, which in a preferred embodiment is the unigram probability of that source morpheme. These unigram probabilities are estimated by transducing a large body of source text and simply counting the frequency with which the different source morphemes occur. The normalizer for a priority queue is defined to be the product of all the unigram probabilities for the morphemes at the positions in the associated set of source structure positions. For example, the normalizer for the priority queue associated with the set [2,3] for the source structure la jeune fille V.sub.— past.sub.— 3s reveiller sa mere is:

Detailed Description Text - DETX (933):

For each priority queue Q on the frontier define the normed score of Q to be equal to the score of the partial hypothesis with the greatest score in Q divided by the normalizer for Q. Let Z be equal to the maximum of all normed scores for all priority queues on the frontier. The threshold assigned to a priority queue Q on the frontier is then equal to Z times the normalizer for that priority queue divided by a constant which in a preferred embodiment is 45.

Detailed Description Text - DETX (934):

After step 6602, thresholds have been assigned to the priority queues on the frontier, a loop is performed in steps 6604 through 6610. The loop count i is equal to a different cardinality on each iteration of the loop. The counter i is initialized in step 6604 to the largest cardinality of any active priority queue, in other words, i is initialized to the maximum cardinality of any priority queue on the frontier. On each iteration of the loop the value of i is decreased by 1 until i is equal to 0, at which point the test 6604 is satisfied and the process of selecting partial hypotheses to be extended is terminated.

Detailed Description Text - DETX (936):

A schematic flow diagram for this processing step 6608 is shown in FIG. 67. The priority queue Q to be processed enters this step at 6701. Steps 6704 through 6707 perform a loop through all partial hypotheses i in the priority queue Q which are greater than the threshold associated with Q. At step 6705 the partial hypothesis i is added to the list of partial hypotheses to be extended. At step 6706 i is used to adjust the thresholds of all active priority queues which are parents of Q. These thresholds are then used when priority queues of lower priority are processed in the loop beginning at step 6604 in FIG. 66.

Detailed Description Text - DETX (937):

Each priority queue which is a parent of partial hypothesis i at step 6706 contains partial hypotheses which account for one less source morpheme than the partial hypothesis i does. For example, consider the partial hypothesis depicted in FIG. 59. Suppose this is the partial hypothesis i. The two target morphemes the and girl are aligned with the three source morphemes la, jeune, and fille which are in source structure positions 1, 2, and 3 respectively. This hypothesis i is therefore in the priority queue corresponding to the set [1,2,3]. The priority queues that are parents of this hypothesis correspond to the sets [1,2], [1,3], and [2,3]. We can use partial hypothesis i to adjust the threshold in each of these priority queues, assuming they are all active, by computing a parent score, score.sub.p from the score score.sub.i associated with the partial hypothesis i. A potentially different parent score is computed for each active parent priority queue. That parent score is then divided by a constant, which in a preferred embodiment is equal to 45. The new threshold for that queue is then set to the minimum of the previous threshold and that parent score.

Detailed Description Text - DETX (938):

These parent scores are computed by removing from score.sub.1 the contributions for each of the source morphemes la, jeune, and fille. For example, to adjust the threshold for the priority queue [2,3], it is necessary to remove the contribution to the score.sub.i associated with the source morpheme in position 1, which is la. This morpheme is the only morpheme aligned with the, so the language model contribution for the must be removed, as well as the translation model contributions associated with la. Therefore: ##EQU110##

Detailed Description Text - DETX (939):

As another example, to adjust the threshold for the priority queue [1,3], it is necessary to remove the contribution to the score.sub.i associated with the source morpheme in position 2, which is jeune. This morpheme is one of two aligned with the target morpheme girl. If the connection between girl and jeune is removed from the partial alignment in FIG. 59, there is still a connection between girl and fille. In other words, girl is still needed in the partial hypothesis to account for fille. Therefore, no language model component is removed. The parent score in this case is: ##EQU111## Here, the first quotient adjust the fertility score, the second adjusts the lexical score and the third adjusts the distortion score.

Detailed Description Text - DETX (940):

With some thought, it will be clear to one skilled in the art how to generalize from these examples to other situations. In general, a parent score is computed by removing a connection from the partial alignment associated with the partial hypothesis i. Such a connection connects a target morpheme t in the partial target structure associated with the partial hypothesis i and a source morpheme s in a source structure. If this connection is the only connection to the target morpheme t, then the language model score for t is divided out, otherwise it is left in. The lexical and distortion scores associated with the source morpheme s are always divided out, as is the fertility score associated with the target morpheme t. If n connections remain to the target morpheme t, since n+1 source morphemes are aligned with t in the partial hypothesis i, then the open fertility score serving as an estimate of the probability that at least n+1 source morphemes will be aligned with t is multiplied in.

Detailed Description Text - DETX (944):

Extensions of partial hypotheses h which are closed are made in lines 17 through 29. First, in line 17 the variable s is set to the identity of the source morpheme at position p in the source structure. This morpheme will have a number of possible target translations. In terms of the translation model, this means that there will be a number of target morphemes t for which the lexical parameter $t(t.verline.s)$ is greater than a certain threshold, which in an embodiment is set equal to 0.001. The list of such target morphemes for a given source morpheme s can be precomputed. In lines 18 through 29 a loop is made through a list of the target morphemes for the source morpheme s. The variable t is set to the target morpheme being processed in the loop. On line 20, an extension is made in which the target morpheme t is appended to the right end of the partial target structure associated with h and then aligned with the source morpheme at position p, and in which the target morpheme t is open in the resultant partial hypothesis. On line 21, an extension is made in which the target morpheme t is appended to the right end of the partial target structure associated with h and then aligned with the source morpheme at position p, and in which the target morpheme t is closed in the resultant partial hypothesis. On line 22, an extension is made in which the target morpheme t is appended to the null target morpheme in the partial target structure associated with hypothesis h. It is assumed throughout this description of hypothesis search that every partial hypothesis comprises a single null target morpheme.

Detailed Description Text - DETX (945):

The remaining types of extensions to be performed are those in which the target structure is extended by two morphemes. In such extensions, the source morpheme at position p is aligned with the second of these two target morphemes. On line 23, a procedure is called which creates a list of target morphemes that can be inserted between the last morpheme on the right of the hypothesis h and the hypothesized target morpheme, t. The lists of target morphemes created by this procedure can be precomputed from language model parameters. In particular, suppose t.sub.tr is the last morpheme on the right of the partial target structure comprised by the partial hypothesis h. For any target morpheme t.sub.1 the language model provides a score for the three-word sequence t.sub.r t.sub.1 t. In one preferred embodiment this score is equal to an estimate of 1-gram probability for the morpheme t.sub.r, multiplied by an estimate of the probability with 2-gram conditional probability with which t.sub.1 follows t.sub.r, multiplied by an estimate of the 3-gram conditional probability with which t.sub.1 follows the pair t.sub.r t.sub.1. By computing such a score for each target morpheme t.sub.1, the target morphemes can be ordered according to these scores. The list returned by the procedure called on line 23 is comprised of the m best t.sub.1 -s which have scores greater than a threshold z. In one embodiment, z is equal to 0.001 and m is equal to 100.

Detailed Description Paragraph Footnote - DEFN (3):

.sup.4 In these equations and in the remainder of the paper, the generic symbol ##EQU94## is used to denote a normalizing factor that converts counts to probabilities. The actual value of ##EQU95## will be implicit from the

context. Thus, for example, in the left hand equation of (168), the normalizing factor is $\text{norm} = \frac{\text{SIGMA}_{f,e} c(f, \text{vertline}, e)}{\text{SIGMA}_{f,e} c(f, \text{vertline}, e)}$ which equals the average length of source sentences. In the right hand equation of (168), the normalizing factor is the average length of target sentences.

Detailed Description Paragraph Table - DETL (37):

TABLE 8 Translation and fertility probabilities for nodding.
 nodding f t(.function..vertline.e) .phi. n(.phi..vertline.e) signe
 0.164 4 0.342 la 0.123 3 0.293 tete 0.097 2 0.167 oui 0.086 1 0.163 fait 0.073 0 0.023 que 0.073 hoche 0.054
 hocher 0.048 faire 0.030 me 0.024 approuve 0.019 qui 0.019 un 0.012 faites 0.011

Detailed Description Paragraph Table - DETL (38):

TABLE 9 Translation and fertility probabilities for should.
 should f t(.function..vertline.e) .phi. n(.phi..vertline.e) devrait
 0.330 1 0.649 devraient 0.123 0 0.336 devons 0.109 2 0.014 faudrait 0.073 faut 0.058 doit 0.058 aurait 0.041
 doivent 0.024 devons 0.017 devrais 0.013

Detailed Description Paragraph Table - DETL (39):

TABLE 10 Translation and fertility probabilities for national.
 national f t(.function..e) .phi. n(.phi..vertline.e) nationale 0.469 1
 0.905 national 0.418 0 0.094 nationaux 0.054 nationales 0.029

Detailed Description Paragraph Table - DETL (40):

TABLE 11 Translation and fertility probabilities for the. the f t
 (.function..vertline.e) .phi. n(.phi..vertline.e) le 0.497 1 0.746 la
 0.207 0 0.254 les 0.155 l' 0.086 ce 0.018 cette 0.011

Detailed Description Paragraph Table - DETL (41):

TABLE 12 Translation and fertility probabilities for farmers.
 farmers f t(.function..vertline.e) .phi. n(.phi..vertline.e) agriculteurs
 0.442 2 0.731 les 0.418 1 0.228 cultivateurs 0.046 0 0.039 producteurs 0.021

Detailed Description Paragraph Table - DETL (42):

TABLE 13 Translation and fertility probabilities for external.
 external f t(.function..vertline.e) .phi. n(.phi..vertline.e) exterieures
 0.944 1 0.967 exterieur 0.015 0 0.028 externe 0.011 exterieurs 0.010

Detailed Description Paragraph Table - DETL (43):

TABLE 14 Translation and fertility probabilities for answer.
 answer f t(.function..vertline.e) .phi. n(.phi..vertline.e) reponse
 0.442 1 0.809 repondre 0.223 2 0.115 repondu 0.041 0 0.074 a 0.038 solution 0.027 repondez 0.021 repondrai
 0.016 reponde 0.014 y 0.013 ma 0.010

Detailed Description Paragraph Table - DETL (44):

TABLE 15 Translation and fertility probabilities for oil. oil f t
 (.function..vertline.e) .phi. n(.phi..vertline.e) petrole 0.558 1 0.760

petrolieres 0.138 0 0.181 petroliere 0.109 2 0.057 le 0.054 petrolier 0.030 petroliers 0.024 huile 0.020 Oil 0.013

Detailed Description Paragraph Table - DETL (45):

TABLE 16 Translation and fertility probabilities for former.
former f t(.function..vertline.e) .phi. n(.phi..vertline.e) ancien 0.592
1 0.866 anciens 0.092 0 0.074 ex 0.092 2 0.060 precedent 0.054 l' 0.043 ancienne 0.018 ete 0.013

Detailed Description Paragraph Table - DETL (46):

TABLE 17 Translation and fertility probabilities for not. not f t
(.function..vertline.e) .phi. n(.phi..vertline.e) ne 0.497 2 0.735 pas
0.442 0 0.154 non 0.029 1 0.107 rien 0.011

Detailed Description Paragraph Table - DETL (47):

English vocabulary e English word e English sentence l length
of e i position in e, $i = 0, 1, \dots, l$ e.sub.i word i of e e.sub.0 the empty notion e.sub.1.sup.i e.sub.1 e.sub.2 . . .
e.sub.i French vocabulary .function. French word f French sentence m length of f j position in f, $j = 1, 2, \dots, m$
m.function..sub.j word j of f .function..sub.1.sup.j .function..sub.1 .function..sub.2function..sub.j a alignment
a.sub.j position in e connected to position j of f for alignment a a.sub.1.sup.j a.sub.1 a.sub.2 . . . a.sub.j .phi..sub.i
number of positions of f connected to position i of e .phi..sub.1.sup.i .phi..sub.1 .phi..sub.2phi..sub.i .tau.
tableau - a sequence of tablets, where a tablet is a sequence of French word .tau..sub.i tablet i of .tau.
.tau..sub.0.sup.i .tau. .sub.0 .tau..sub.1tau..sub.i .phi..sub.i length of .tau..sub.i k position within a tablet, $k = 1, 2, \dots, \text{phi..sub.i}$.tau..sub.ik word k of .tau..sub.i .pi. a permutation of the positions of a tableau .pi..sub.ik
position in f for word k of .tau..sub.i for permuta- tion .pi. .pi..sub.i1.sup.k .pi..sub.i1 .pi..sub.i2pi..sub.ik V
(f.vertline.e) Viterbi alignment for (f.vertline.e) V.sub.i.rarw.j (f.vertline.e) Viterbi alignment for (f.vertline.e) with ij
pegged (a) neighboring alignments of a .sub.ij (a) neighboring alignments of a with ij pegged b(a) alignment in (a)
with greatest proba- bility b.sub.infin. (a) alignment obtained by applying b repeatedly to a b.sub.i.rarw.j (a)
alignment in .sub.ij (a) with greatest probability b.sub.i.rarw.j.sup.infin. (a) alignment obtained by applying
b.sub.i.rarw.j repeatedly to a (e) class of English word e (.function.) class of French word .function. .DELTA.j
displacement of a word in f .nu., .nu.' vacancies in f .rho..sub.i first position in e to the left of i that has non-zero
fertility c.sub.i average position in f of the words connected to position i of e [i] position in e of the i.sup.th one
word notion .sub.i c.sub.[i] P.sub.theta. translation model P with parameter values .theta. C(f, e) empirical
distribution of a sample .psi.(P.sub..theta.) log-likelihood objective function R(P.sub.theta., P.sub.theta.) relative
objective function t(f.vertline.e) translation probabilities (All Models) .epsilon.(m.vertline.l) sentence length
probabilities (Models 1 and 2) n(.phi..vertline.e) fertility probabilities (Models 3, 4, and 5) .rho..sub.0, .rho..sub.1
fertility probabilities for e.sub.0 (Models 3, 4, and 5) a(i.vertline.j, l, m) alignment probabilities (Model 2) d
(j.vertline.i, l, m) distortion probabilities (Model 3) d.sub.1 (.DELTA.j.vertline. ,) distortion probabilities for the first
word of a tablet (Model 4) d.sub.>1 (.DELTA.j.vertline.) distortion probabilities for the other words of a tablet
(Model 4) d.sub.LR (l.sub.- or sub.- r.vertline.B, .nu., .nu.) distortion probabilities for choosing left or right
movement (Model 5) d.sub.L (.DELTA.j.vertline. , .nu.) distortion probabilities for leftward movement of the first
word of a tablet (Model 5) d.sub.R (.DELTA.j.vertline. , .nu.) distortion probabilities for rightward movement of the
first word of a tablet (Model 5) d.sub.>1 (.DELTA.j.vertline. , .nu.) distortion probabilities for movement of the other
words of a tablet (Model 5)

Claims Text - CLTX (3):

generating the target text in the second language based on a combination of a probability of occurrence of an intermediate structure of text associated with a target hypothesis selected from the second language using a target language model, and a probability of occurrence of the source text given the occurrence of said intermediate structure of text associated with said target hypothesis using a target-to-source translation model; and

Claims Text - CLTX (10):

estimating, for each target hypothesis, a first probability of occurrence of said text associated with said target hypothesis using a target language model;

Claims Text - CLTX (11):

estimating, for each target hypothesis, a second probability of occurrence of the source text given the occurrence of said text associated with said target hypothesis using a target-to-source translation model;

Claims Text - CLTX (12):

combining, for each target hypothesis, said first and second probabilities to produce a target hypothesis match score; and

Claims Text - CLTX (24):

determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to a probability that a notion in the target text will generate a specific number of units of linguistic structure in the source text;

Claims Text - CLTX (25):

determining a alignment score for each of said partial hypotheses, said alignment score being proportional to a probability that a unit of linguistic structure in the target text will align with one of zero or more units of linguistic structure in the source text;

Claims Text - CLTX (26):

determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to the probability that said units of linguistic structure in the target text of a given partial hypothesis will translate into said units of linguistic structure of said source text;

Claims Text - CLTX (27):

determining a distortion model score for each of said partial hypotheses, said distortion model score being proportional to the probability that a source unit of linguistic structure will be in a particular position given a position of the target units of linguistic structure that generated it; and

Claims Text - CLTX (30):

determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to a probability that a notion in the target text will generate a specific number of notion units of linguistic structure in an intermediate structure source text;

Claims Text - CLTX (31):

determining an alignment score for each of said partial hypotheses, said alignment score being proportional to the probability that a unit of linguistic structure in said intermediate structure of target text will align with one of zero or more units of linguistic structure in said intermediate structure of source text;

Claims Text - CLTX (32):

determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to a probability that said units of linguistic structure in said intermediate structure of target text of a given partial hypothesis will translate into said units of linguistic structure of said intermediate structure source text;

Claims Text - CLTX (33):

determining a distortion model score for each of said partial hypotheses, said distortion model score being proportional to a probability that a source unit of linguistic structure will be in a particular position given the position of the target units of linguistic structure that generated it; and

Claims Text - CLTX (43):

estimating a first score, said first score being proportional to a probability of occurrence of each intermediate target structure of text associated with said target hypotheses using a target structure language model;

Claims Text - CLTX (44):

estimating a second score, said second score being proportional to a probability that said intermediate target structure of text associated with said target hypotheses will translate into said intermediate source structure of text using a target structure-to-source structure translation model;

Claims Text - CLTX (79):

21. A method according to claim 10, wherein each of said intermediate target structures is expressed as an ordered sequence of units of linguistic structure, and a first score probability is obtained by multiplying conditional probabilities of each of said units of linguistic structure within an intermediate target structure given an occurrence of previous units of linguistic structure within said intermediate target structure.

Claims Text - CLTX (80):

22. A method according to claim 21, wherein said conditional probability of each unit of each of linguistic structure within an intermediate target structures depends only on a fixed number of preceding units within said intermediate target structure.

Claims Text - CLTX (84):

means for estimating, for each target hypothesis, a first probability of occurrence of said text associated with said target hypothesis using a target language model;

Claims Text - CLTX (85):

means for estimating, for each target hypothesis, a second probability of occurrence of the source text given the occurrence of said text associated with said target hypothesis using a target-to-source translation model;

Claims Text - CLTX (86):

means for combining, for each target hypothesis, said first and second probabilities to produce a target hypothesis match score; and

Claims Text - CLTX (98):

means for determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to the probability that a notion in the target text will generate a specific number of notion units of linguistic structure in an intermediate structure source text;

Claims Text - CLTX (99):

means for determining an alignment score for each of said partial hypotheses, said alignment score being

proportional to a probability that a unit of linguistic structure in said intermediate structure of target text will align with one of zero or more units of linguistic structure in said intermediate structure of source text;

Claims Text - CLTX (100):

means for determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to the probability that said units of linguistic structure in said intermediate structure of target text of a given partial hypothesis will translate into said units of linguistic structure of said intermediate structure source text;

Claims Text - CLTX (101):

means for determining a distortion model score for each of said partial hypotheses, said distortion model score being proportional to a probability that a source unit of linguistic structure will be in a particular position given the position of the target units of linguistic structure that generated it; and

Claims Text - CLTX (104):

means for determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to the probability that a notion in the target text will generate a specific number of notions of units of linguistic structure in the source text;

Claims Text - CLTX (105):

means for determining an alignment score for each of said partial hypotheses, said alignment score being proportional to the probability that a unit of linguistic structure in the target text will align with one of zero or more units of linguistic structure in the source text;

Claims Text - CLTX (106):

means for determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to the probability that said units of linguistic structure in the target text of a given partial hypothesis will translate into said units of linguistic structure of source text;

Claims Text - CLTX (107):

means for determining a distortion model score for each of said partial hypothesis, said distortion model score being proportional to the probability that the source units of linguistic structure will be in a particular position given the position of the target units of linguistic structure that generated it; and

Claims Text - CLTX (118):

means for estimating a first score, said first score being proportional to a probability of occurrence of each intermediate target structure of text associated with said target hypotheses using a target structure language model;

Claims Text - CLTX (119):

means for estimating a second score, said second score being proportional to a probability that said intermediate target structure of text associated with said target hypotheses will translate into said intermediate source structure of text using a target structure-to-source structure translation model;

Claims Text - CLTX (152):

means for expressing each of said intermediate target structures as an ordered sequence of units of linguistic structure, and means for multiplying conditional probabilities of said units within an intermediate target structure

given an occurrence of previous units of linguistic structure within said intermediate target structure to obtain said first score.

Claims Text - CLTX (153):

42. A system according to claim 41, wherein a conditional probability of each unit of each of said intermediate target structure depends only on a fixed number of preceding units within said intermediate target structure.

Other Reference Publication - OREF (8):

"Interpolated Estimation of Markov Source Parameters from Sparse Data" by F. Jelinek and R. L. Mercer, Workshop on Pattern Recognition in Practice, Amsterdam (Netherland): North Holland, May 21-23, 1980.

Other Reference Publication - OREF (9):

"An Inequality and Associated Maximization Technique in Statistical Estimation for Probabilistic Functions of Markov Processes", by Leonard E. Baum, Inequalities, vol. 3, 1972, pp. 1-8.

Other Reference Publication - OREF (12):

"Deriving Translation Data from Bilingual Texts", Catizone et al., Proceedings of the First International Acquisition Workshop, Detroit, Mich., 1989.